

DIP Protocol

The Decentralised Intellectual Property

Peeramid Labs
The AI Futures Collective
T. Pecerskis, A. Smirnovs, et al.
Email: contact@peeramid.xyz

Abstract—The proliferation of advanced Artificial Intelligence (AI) has ignited an unprecedented demand for high-quality, verifiable data, yet the digital economy lacks a robust framework for valuing and protecting this foundational intellectual property (IP). Simultaneously, the promise of decentralization has been hampered by the geographic centralization of its node infrastructure and governance models vulnerable to capture. This paper introduces the Decentralised Intellectual Property (DIP) Protocol, a novel framework designed to solve these parallel challenges. DIP establishes a secure and liquid market for digital IP by integrating four core innovations: a meritocratic system for generating data with verifiable provenance; a novel economic architecture that enables Meritocratic Autonomous Organizations (MAOs) where governance is dynamic and competitive; a flexible infrastructure for privacy-preserving computation; and a model for incentivizing true geospatial decentralization. By enabling parties to verify the properties of a dataset without requiring access to the data itself, and by creating a framework for unified, user-centric digital marketplaces, DIP provides the foundational layer for a more equitable, efficient, and institutionally-viable ownership economy.

Index Terms—Decentralized Intellectual Property; Data Markets; Zero-Knowledge Proofs; ZKML; Multi-Party Computation; DAO; Tokenomics; AI; Machine Learning.

I. INTRODUCTION: THE EMERGING DATA ECONOMY AND THE INTELLECTUAL PROPERTY PARADOX

A. The AI/ML Revolution and the Insatiable Demand for Verifiable Data

The contemporary technological landscape is being fundamentally reshaped by advancements in Artificial Intelligence (AI) and Machine Learning (ML). The capabilities of these systems, from large language models to sophisticated predictive analytics, are directly contingent upon the volume, relevance, and, most critically, the quality of the data on which they are trained. The foundational principle of "Garbage In, Garbage Out" (GIGO) dictates that even the most advanced algorithms will produce inaccurate, biased, or unreliable outcomes if trained on flawed data. This dependency has created an insatiable demand not merely for data, but for high-fidelity, verifiable, and "fit for purpose" datasets [1], [2].

The economic and strategic risks associated with poor data quality are substantial. They manifest as inaccurate predictions, the amplification of societal biases, and significantly inflated project costs, as data scientists expend a majority of their efforts on data cleaning and preparation rather than model development [3]. With AI adoption accelerating—78% of organizations reported using AI in at least one business function

in 2024, supported by record levels of private investment—the market for premium data is expanding at an unprecedented rate [4]–[6].

This explosion in AI has also surfaced a new, systemic threat: model collapse. Recent research indicates that generative models trained recursively on synthetic or AI-generated data risk degrading in quality, losing diversity, and amplifying biases over time [7]. A study from Meta highlighted that even 1% of synthetic data in a training set can lead to this degenerative phenomenon, where larger datasets no longer yield performance improvements. Similarly, research from Anthropic has demonstrated a phenomenon termed "subliminal learning," where AI models can covertly transmit hidden behaviors and biases to other models through seemingly random data, a risk that is magnified by the industry's increasing reliance on synthetic data to reduce costs [8]. These challenges underscore a critical market need: a reliable and scalable source of high-quality, verifiably human-generated data to serve as a "ground truth" for training, fine-tuning, and re-calibrating the next generation of AI systems.

B. The Web3 Dilemma: Open-Source Ideals vs. The Unsolved IP Problem

The Web3 ecosystem, built on the principles of decentralization and permissionless innovation, has operated primarily under an open-source ethos. While this has been a powerful catalyst for development, it has simultaneously created a structural inability to value, protect, and monetize intellectual property (IP). The prevailing paradigm encourages the creation of digital public goods but offers no native mechanism to compensate their creators, leading to an extractive dynamic where centralized commercial entities can freely capitalize on open-source code and public data without remuneration [7]. This represents a significant market failure; for Web3 to mature and capture value at the scale of the traditional economy, it must evolve beyond a purely open-source model and develop robust frameworks for managing IP [9].

The challenges of protecting IP in a decentralized environment are profound. Traditional legal frameworks rely on centralized authorities for enforcement, a model that is incompatible with the pseudonymous, borderless, and immutable nature of blockchains. Issues of anonymity, jurisdictional ambiguity, and the finality of on-chain transactions render conventional IP enforcement mechanisms largely ineffective [10]–[12]. The market has begun to recognize the immense value in solving

these deep, structural problems. The significant valuation of companies like Zama, which focuses on providing on-chain confidentiality through advanced cryptography, serves as a powerful market signal that verifiable privacy and data control are not peripheral features but core components of Web3's future value proposition [13]. The open-source ethos of Web3 can be understood not just as a philosophical choice, but as a direct technical consequence of its failure to solve the underlying data control problem. In the absence of a trusted central authority, verifying the value of a digital asset has historically required direct access to it. However, in a trustless environment, granting access is tantamount to relinquishing control. Faced with this paradox, the only viable path to permissionless collaboration was to make the IP—the code and the data—open by default. A technical solution that allows for verification without disclosure would fundamentally break this dependency, enabling a new paradigm for proprietary IP in decentralized ecosystems.

C. The Data Disclosure Challenge: Why Traditional IP Fails in Trustless Environments

The core technical barrier preventing the emergence of a robust market for digital IP is the data disclosure problem. This refers to the fundamental paradox that to prove the value or verify the properties of a piece of digital information, one has historically been required to disclose the information itself, thereby destroying its scarcity and commercial value as exclusive IP [14], [15]. Any party who is granted a copy of a dataset for evaluation can subsequently redistribute that copy publicly, making the asset worthless. This problem has long been considered intractable in trustless systems, confining valuable data to centralized, walled-garden platforms where trust is enforced through legal agreements and platform-level controls rather than technical guarantees [16], [17].

This challenge is precisely why a true, decentralized data economy has failed to materialize. Without a cryptographic solution, the exchange of valuable data requires trusting a counterparty or an intermediary, reintroducing the very centralization that Web3 seeks to eliminate. The problem is not merely one of preventing unauthorized copying (a task at which digital rights management has largely failed) but of enabling confident commercial exchange. A buyer needs to verify that a dataset meets their specifications before purchase, but the seller cannot provide the data for verification without losing control over it. This dilemma has been the primary impediment to the creation of liquid, open markets for proprietary data. The solution to this long-standing problem lies in advanced cryptographic primitives that can decouple verification from disclosure, allowing for computation and validation to be performed on encrypted or otherwise obscured data.

D. Thesis: Introducing DIP as a Foundational Layer for Verifiable Digital IP

This paper introduces the Decentralised Intellectual Property (DIP) Protocol, a novel framework designed to create a

secure, transparent, and liquid market for digital IP by directly solving the data disclosure problem and the challenge of verifiable network decentralization. The protocol's architecture is built upon four integrated pillars:

- 1) **Meritocratic Provenance:** A system for generating data with a verifiable and auditable history, rooted in the demonstrated competence of its human and AI contributors. This ensures that data quality is not a subjective label but an emergent property of a rigorous, incentive-aligned process.
- 2) **Meritocratic Autonomous Organizations (MAOs):** A sophisticated economic model that evolves the DAO concept, enabling the creation of unified digital marketplaces governed by dynamic, competitive, and merit-based principles rather than static token-based voting.
- 3) **Privacy-Preserving Infrastructure:** A flexible, L2-native network of service nodes utilizing a spectrum of cryptographic techniques, including Multi-Party Computation (MPC) and Zero-Knowledge Machine Learning (ZKML), to solve the data disclosure problem.
- 4) **Geospatial Decentralization:** A node participation model that moves beyond simplistic token staking to incentivize tangible, real-world infrastructure investment, providing a systemic solution for on-chain institutional compliance and resilience.

The demand for high-quality data from the AI industry represents the first "killer application" for such a protocol. Unlike many previous Web3 use cases that created self-contained economies, the AI data market is a massive, external, and established economy that Web3 is currently ill-equipped to service [18]–[20]. By providing the infrastructure for verifiable quality, secure access, and liquid exchange, DIP aims to act as a bridge, enabling the decentralized world to capture a significant share of this multi-trillion dollar market. The following sections will provide a detailed exposition of the protocol's theoretical underpinnings, economic model, and technical infrastructure.

II. A NEW FOUNDATION: MERITOCRATIC PROVENANCE AND VERIFIABLE COMPETENCE

A. Redefining Data Quality: From Subjective Labels to Provable Merit

In the context of the data economy, the term "data quality" is often ambiguous and subjective. The DIP Protocol posits that a more robust and useful concept is data provenance: a cryptographically verifiable and immutable record of a data asset's entire lifecycle, from its origin and creation process to the demonstrated competence of the agents who contributed to it [21], [22]. Within this framework, true quality is not an externally applied label but an emergent and intrinsic property of a well-defined, meritocratic, and transparent generation process. The protocol reframes the central question from "How can we verify that this data is high-quality?" to "How can we design a system that reliably produces high-quality data as its natural output?" The answer lies in establishing a

verifiable measure of competence for the agents—both human and AI—who create and curate the data.

B. The Autonomous Competence Identification Protocol (ACIP) as the Engine of Provenance

The core mechanism for establishing data provenance in the DIP Protocol is the Autonomous Competence Identification Protocol (ACIP), a novel system for ranking agents in trustless environments based on verifiable merit [7], [23], [24]. ACIP provides the foundational layer for ensuring that data is generated by agents with demonstrated expertise in a given domain.

ACIP formally defines competence as a quantifiable measure derived from the time and financial resources an agent is willing to commit to a specific domain, validated by the peer-assessed recognition they receive from fellow participants [7]. This creates a dynamic ranking ladder where participants join tiered groups and engage in competitive interactions, or “elections,” to demonstrate their skills. Advancement to higher ranks requires succeeding in these time-locked, cost-based competitions.

The protocol’s design is inherently resistant to Sybil attacks. The cost to participate in a ranking event, X_{id} , is a function of the time commitment, T_{id} , and protocol-wide constants. To deterministically manipulate the system and achieve a high rank, an adversary would need to win multiple sequential ranking events. The total cost to achieve Rank R , $C_{TotalSybilToRankR}$, scales exponentially with the rank, as described by the equation (Eq. 1).

where N_{min} is the minimum number of participants per group and $X_{avg_group_cost}$ is the average cost to control a single group instance [7]. This exponential cost structure makes sustained, high-rank Sybil attacks prohibitively expensive for a rational actor.

Furthermore, ACIP moves beyond a single, monolithic score by representing competence as a multi-dimensional vector, $P_r = (c_1, c_2, \dots, c_m)$, where each component quantifies proficiency in a distinct domain (e.g., image annotation, code analysis, legal review). This allows for a nuanced and verifiable representation of an agent’s specific expertise, forming the basis for a rich “competence multiverse” [7].

C. Collaborative Curation: Generating High-Fidelity Datasets with the Continuous Voting-Proposing Protocol (CVPP)

While ACIP provides the framework for identifying competent agents, the Continuous Voting-Proposing Protocol (CVPP) provides a practical methodology for these agents to engage in collaborative knowledge work, such as the creation of high-fidelity datasets [7], [25]. CVPP is a structured, gamified protocol that facilitates group consensus and ordering of information through iterative rounds of private proposals and voting.

The protocol operates in distinct stages:

- 1) **Proposing Stage:** Participants privately submit proposals (e.g., data points, labels, or other contributions) to a neutral facilitator or smart contract.

- 2) **Voting Stage:** The aggregated, anonymized proposals are presented to the group, and participants vote on their preferred submissions.
- 3) **Reveal Stage:** Scores are tallied, and the results, including proposer identities, are revealed.

This process, particularly the private nature of proposing and voting, is designed to mitigate common cognitive biases such as the halo effect or bandwagoning, encouraging participants to evaluate contributions on their intrinsic merit [7]. Empirical studies of a CVPP implementation for collaborative playlist curation demonstrated high participation rates (95% for active users) and high user satisfaction with the quality of the output (92% rating), validating its efficacy as a mechanism for generating community-vetted, high-quality ordered lists, which serve as a direct proxy for curated datasets [7].

The integration of ACIP and CVPP creates a powerful data generation engine. ACIP identifies and assembles groups of verifiably competent experts for a specific task, and CVPP provides them with a structured and incentive-aligned process to collaboratively produce a high-quality dataset. The entire process—the proposals, the votes, the evolution of consensus—is recorded on-chain, creating a rich, auditable trail of the dataset’s provenance. This trail is the dataset’s metadata, embedded at the moment of creation.

D. From Competence to Quality: A Formal Model ($q = f(x, t)$)

The synthesis of the ACIP and CVPP frameworks allows for a formal, quantifiable model of data quality. We can define the quality of a dataset, q , as a function of the energy and time invested in its creation:

$$q = f(x, t) \quad (2)$$

In this model:

- q represents the data quality, a verifiable and predictable output of the protocol. It is a measure of the dataset’s fidelity, accuracy, and relevance, as determined by a meritocratic peer-review process.
- x represents the energy committed to the process. This is the irreversible financial stake, X_{id} , required for participation in an ACIP ranking group, which serves as a costly signal of commitment and a deterrent to low-effort or malicious contributions [7].
- t represents the time commitment, T_{id} , required to finalize a competitive interaction within a group. This ensures sustained engagement and prevents superficial participation [7].
- f represents the process function itself—the set of rules and interactions defined by the ACIP and CVPP protocols. It is the competitive, peer-validated mechanism that transforms the inputs of time and energy into the output of verifiable quality.

The core assertion of the DIP Protocol is that by precisely defining and controlling the inputs (x, t) and the process (f), the system can reliably produce data of a predictable and

$$C_{TotalSybilToRankR} \approx \sum_{k=0}^{R-1} N_{min}^{R-1-k} \cdot X_{avg_group_cost} = X_{avg_group_cost} \cdot \frac{N_{min}^R - 1}{N_{min} - 1} \quad (1)$$

verifiable quality level, q . This allows the protocol to programmatically target and generate datasets that meet specific quality thresholds, such as the "top quartile" of all datasets within a given domain, and to provide cryptographic assurances of this quality to consumers. This model of meritocratic data generation offers a systemic defense against the threat of AI model collapse. By requiring verifiable commitments of time and economic stake, the protocol creates a high barrier to entry that is straightforward for motivated humans to meet but economically infeasible for automated agents to fake at scale. The output is a continuous stream of data with strong cryptographic and economic guarantees of its human-centric, high-effort origin. This stream can serve as a trusted "gold standard" for the AI industry, providing a crucial and verifiable source of ground-truth data to benchmark, fine-tune, and safeguard AI models against degradation.

III. ECONOMIC ARCHITECTURE: LIQUID ACCESS AND ALIGNED INCENTIVES

A. Liquid Access Tokens (LATs): A Novel Primitive for Quantized Data Quality

The core economic primitive of the DIP Protocol is the **Liquid Access Token (LAT)**. A LAT is a fungible (ERC-20) or semi-fungible (ERC-1155) digital asset that represents a tokenized right to access a specific dataset or data stream at a defined quality level, q . Each LAT is a claim on a piece of intellectual property whose provenance and quality have been established through the ACIP and CVPP mechanisms.

The defining innovation of LATs is their liquidity. Unlike traditional, static access credentials or the non-transferable Soulbound Tokens (SBTs) proposed in other identity systems, LATs are designed to be freely tradable, creating a dynamic, open market for data access [7]. This transforms intellectual property from an illiquid, difficult-to-price asset into a standardized, composable, and liquid financial instrument.

The tokenomics of LATs are flexible and can be tailored to specific use cases, enabling a variety of access models [26]–[28]:

- **One-Time Access:** A LAT can be designed to be burned upon use, granting a single, exclusive access to the underlying data. This model is ideal for buyers who require data exclusivity for training a proprietary AI model.
- **Subscription Access:** Holding a certain number of LATs could grant continuous access to a data stream for a defined period, functioning like a decentralized subscription service.
- **Tiered Access:** The quantity of LATs held could determine the level or granularity of data access, allowing for more sophisticated pricing and packaging of IP.

By quantizing data quality (q) and representing it as a liquid token, the DIP Protocol creates the necessary conditions for

efficient price discovery and allocation of intellectual property in a decentralized market.

B. The Three-Party Escrow Model: An Advanced Use Case for High-Value IP

While the DIP protocol provides a flexible foundation for various data commerce models, its full power is demonstrated in a novel three-party escrow system designed for high-value, enterprise-grade IP exchange. This model is designed based on game-theoretic principles to align the incentives of three distinct, rational agents: the Buyer, the Expert Community, and the Expert Guild [29], [30].

- **The Buyer (Purchaser):** An entity (e.g., an AI company) seeking to acquire high-quality data. The Buyer initiates a transaction by issuing a Request-for-Data (RFD). This RFD is a smart contract that specifies the desired data type, the required quality level (q), and any other parameters (such as the ZKML verification model, detailed in Section 4). The Buyer funds this contract with the payment, which is held in escrow. Their primary objective is to acquire verifiably clean data with minimal counterparty risk.
- **The Expert Community (Workers):** A decentralized network of individuals or AI agents who participate in the ACIP/CVPP framework to generate data. By contributing their time and expertise, they earn LATs corresponding to the quality of the data they produce. Their primary objective is to monetize their skills by selling their earned LATs to fulfill a Buyer's RFD.
- **The Expert Guild (DAO as Steward):** A Decentralized Autonomous Organization (DAO) that serves as the trusted third-party escrow agent, arbitrator, and long-term *steward* of the intellectual property. Guilds are formed by high-ranking experts who have achieved a significant level of competence within a specific domain and have chosen to transition from active data creation to a governance and curation role. The Guild's primary objective is to maintain the long-term value of its ecosystem. This includes not only facilitating transactions but also ensuring the ongoing availability and monetization of the data assets created by its community. As such, the Guild is the recipient of **Data Retention Fees** from Buyers, and it is the Guild's responsibility to use these fees to contract with and incentivize DIP Nodes for secure, long-term data hosting. This aligns the Guild's incentives with the perpetual value of the IP, rather than just the initial sale.

This three-party model creates a system of checks and balances where each participant is incentivized to act honestly to maximize their own utility, leading to a stable and self-regulating market equilibrium.




Fig. 1. The Three-Party Escrow Model. The Buyer issues an RFD, Workers provide data for LATs, and the Guild acts as a reputation-based arbiter, creating a game-theoretically balanced system.

C. Tokenomics of LATs: Governance, Staking, and Liquidation Dynamics

The incentive alignment at the heart of the escrow model is driven by a unique tokenomic mechanism centered on the relationship between LATs and the governance of the Expert Guild.

A fundamental feature of the protocol is that experts who earn LATs have the option to convert them into the Guild’s native governance tokens. This provides a valuable “exit” opportunity for experts, allowing them to transition from being active workers to becoming stakeholders in the governance and long-term success of the ecosystem. However, this conversion right creates a central economic tension, which can be termed the **Dilution Dilemma**. From the perspective of the existing governance token holders of the Guild, every LAT conversion represents a dilution of their voting power and their claim on the Guild’s future revenue. Therefore, while the Guild is broadly incentivized to foster a vibrant ecosystem to attract more business, its incumbent governors are specifically incentivized to minimize the rate at which new governance tokens are minted via LAT conversion.

The Buyer’s RFD is strategically designed to leverage this tension. A key condition of the RFD smart contract is that any LATs purchased to fulfill the request must be permanently locked or burned. This clause makes the transaction highly attractive to the Expert Guild. By agreeing to act as the escrow and arbitrator for the RFD, the Guild not only facilitates economic activity (from which it may derive fees) but also ensures that this specific transaction eliminates a quantum of potential governance dilution.

This dynamic establishes a robust game-theoretic equilibrium:

- The Guild is strongly motivated to provide fair and rigorous escrow services to attract more Buyers. Each successful transaction reinforces its reputation and, through the LAT burn mechanism, protects its governance from dilution.
- The Buyer is motivated to use the Guild’s escrow service because the Guild’s self-interest in maintaining its reputation and preventing dilution makes it a credible and aligned arbiter of data quality.
- The Experts are motivated to produce high-quality data

to earn LATs, confident that they have two viable paths to monetization: selling directly to a Buyer’s RFD or converting their LATs into a long-term governance stake.

In the event of a dispute where a Buyer deems the provided data to be inappropriate, the resolution is governed by this incentive structure:

- If the Guild sides with the Buyer against the Workers: The Workers’ staked LATs are slashed or burned, and the payment is returned to the Buyer. This outcome is favorable to the Guild’s governance holders (as it prevents dilution) and reinforces its reputation for upholding high standards with Buyers. However, if this power is used unfairly, it will damage the Guild’s reputation with the Expert Community, potentially causing them to migrate to a competing Guild.
- If the Guild sides with the Workers against the Buyer: The payment is released to the Workers. This maintains the trust of the expert base but risks damaging the Guild’s reputation among Buyers, who may see the Guild as insufficiently rigorous.
- If the Workers and Buyer agree to bypass the Guild: This scenario is disincentivized by the protocol design, as the Guild provides the foundational trust, security, and dispute resolution framework that makes the permissionless transaction viable in the first place.

This system creates a self-regulating market for trust. A Guild that arbitrates poorly will lose market share as both Buyers and Experts migrate to Guilds with a more balanced and reputable track record. This competitive pressure forces Guilds to optimize for long-term fairness and sustainability, effectively creating a decentralized, market-driven judicial system for intellectual property.

D. From DAOs to MAOs: A Paradigm of Competitive Governance

The protocol’s architecture enables a profound evolution of decentralized governance, from the standard Decentralized Autonomous Organization (DAO) to a **Meritocratic Autonomous Organization (MAO)**. In a traditional DAO, a 51% governance capture is often seen as a catastrophic failure. In an MAO, powered by the principles of ACIP, temporary quorum capture is an expected and acceptable part of a dynamic, competitive system.

This paradigm shift is possible because ACIP’s time-financial constraint, $q = f(x, t)$, makes permanent, centralized control economically irrational. Any party wishing to seize governance control faces a clear trade-off: the faster they wish to acquire power (reducing time, t), the exponentially higher the energy and capital cost they must expend (x). This creates a system where governance control is likely to be in a constant state of flux, oscillating between competing parties who are judged by the merit of their decisions, rather than being captured permanently by a single, entrenched entity.

E. Mitigating Collusion and Ensuring Rational Agency

The primary threat to any multi-party system is collusion. The DIP Protocol incorporates several layers of defense, inherited primarily from the design of ACIP, to ensure that collusion between any two parties against the third is economically irrational or technically infeasible over the long term.

The first line of defense is the cost of participation. As established in ACIP, achieving a high rank and the ability to generate valuable LATs requires a significant and irreversible investment of time and capital [7]. This "skin in the game" disincentivizes short-term, extractive behavior.

The second layer of defense is transparency and randomization. All core interactions—group formation, voting, and escrow transactions—are recorded on-chain. This transparency allows for the application of collusion clustering analysis, where voting patterns and social graphs can be monitored to identify and flag coordinated, non-meritocratic behavior [7]. Furthermore, ACIP can incorporate randomized participation in ranking events, making it difficult for colluding groups to guarantee they are placed in the same competitive instance [7].

Finally, the economic model itself provides a strong defense. For example, if a Guild and a group of Workers were to collude to defraud a Buyer with low-quality data, the public and immutable nature of the blockchain record would permanently damage the Guild's reputation. The short-term gain from the single fraudulent transaction would likely be outweighed by the long-term loss of future business from other Buyers, who would shun the now-untrustworthy Guild. This aligns with research on collusion-resistant mechanisms, which demonstrates that well-designed incentive structures can make joint deviation from honest behavior unprofitable [31]–[33]. The protocol is designed such that the dominant strategy for all participants is long-term, honest cooperation.

IV. PROTOCOL INFRASTRUCTURE: VERIFIABLE DECENTRALIZATION AND PRIVACY

A. DIP Nodes: A Foundation for Jurisdictional and Physical Decentralization

The foundational infrastructure of the DIP Protocol is a decentralized network of **DIP Nodes**. These nodes are independent operators who provide the critical computational and storage resources required for the protocol to function. However, unlike validators in traditional Proof-of-Stake networks whose primary contribution is capital-based signing, DIP Nodes provide tangible, verifiable services. Their core responsibilities are:

- **Privacy-Preserving Computation:** Performing authorized encryption, decryption, and participating in the MPC schemes for key management.
- **Provenance Attestation:** Signing and attesting to the integrity of data generation processes.

This service-oriented role provides a systemic solution to the challenge of blockchain decentralization. Because node operators must compete on the quality and security of their services, they are incentivized to make tangible investments

in their infrastructure. This creates a market where nodes can differentiate themselves based on:

- **Jurisdictional Guarantees:** A node can operate within a specific legal jurisdiction (e.g., Switzerland, Singapore) to appeal to institutional clients with strict data residency requirements.
- **Verifiable Security:** Node operators are incentivized to obtain security certifications (e.g., SOC 2, ISO 27001) and even allow physical inspection of their data centers to attract high-value clients.
- **Specialized Hardware:** Nodes may offer specialized hardware, such as Trusted Execution Environments (TEEs) or high-performance GPUs for ZK proof generation, creating a tiered market for different levels of security and performance.
- **Decentralization:** While we speak of DIP node as singular entity, we assume that node can act as aggregated signature, acting as a gatekeeper for larger, underlying network of nodes such that could provide high redundancy and security by own design, examples of such could be Zama, Internet Computer, IEXEC or others.

This transforms decentralization from an abstract metric into a verifiable, market-driven feature, solving a major adoption barrier for enterprise and institutional use of blockchain technology.

1) *A Market for Privacy Technologies:* This market-driven approach also extends to the underlying privacy technologies employed by the nodes. While cutting-edge solutions like Fully Homomorphic Encryption (FHE) offer the ability to compute on encrypted data, the DIP protocol does not mandate its exclusive use. A core design principle is to foster a competitive marketplace where different technological trade-offs can coexist. The security of any advanced cryptographic primitive is strongly dependent on its specific implementation, and for many use cases, the performance overhead of FHE may be unnecessary. Therefore, the protocol allows for a heterogeneous network where different nodes can offer different solutions:

- **FHE-based Nodes:** For high-assurance applications requiring complex computations directly on encrypted data, where performance is a secondary concern to absolute privacy.
- **Hardware-Isolated Nodes:** Utilizing Trusted Execution Environments (TEEs) like Intel SGX or AMD SEV to provide practical privacy with high performance for simpler tasks, appealing to users who accept hardware-based trust models.
- **Standard MPC Nodes:** Focusing on robust key management and secure multi-party computation for access control without relying on specialized hardware or the computational intensity of FHE.

This flexibility allows users and data creators to select the node that provides the optimal balance of security, performance, and cost for their specific needs, rather than enforcing a one-size-fits-all cryptographic solution.

TABLE I
COMPARISON OF IP MANAGEMENT PARADIGMS

Feature	Traditional Legal System	NFT Licensing (e.g., CC0)	DIP Protocol
Proof of Origin	Relies on registration with centralized authorities (e.g., patent offices). Can be slow, expensive, and geographically limited.	On-chain token minting provides a time-stamped proof of creation, but does not verify the creator's identity or originality.	Cryptographically signed actions within ACIP/CVPP create an immutable, on-chain record of contribution tied to a meritocratic identity.
Proof of Quality	Relies on subjective, ex-post evaluation (e.g., market success, expert reviews). No objective, a priori guarantee.	None. The quality of the underlying asset is completely divorced from the NFT that represents it.	Quality (q) is a quantifiable, verifiable output of the meritocratic generation process. Can be objectively proven using ZKML.
Access Control	Enforced through legal agreements and technological restrictions (DRM), which are often brittle and centralized.	Generally none. Access is often public, with the NFT representing a social claim of ownership rather than technical control.	Access is cryptographically enforced via LATs. The protocol allows for granular, programmable access logic (e.g., one-time, subscription).
Liquidity	Highly illiquid. IP transactions are bespoke, high-friction legal processes that can take months or years.	Liquid market for the NFT "wrapper," but not necessarily for the underlying IP rights, which are often ambiguous.	High. LATs are designed as liquid, fungible/semi-fungible assets, creating a dynamic, low-friction market for IP access rights.
Dispute Resolution	Centralized, expensive, and slow court systems. Access to justice can be prohibitive for smaller creators.	Rudimentary and often off-chain. Relies on platform-level moderation or social consensus, with no binding enforcement mechanism.	Decentralized, game-theoretically balanced arbitration provided by Expert Guilds. Incentives are aligned for fast, fair, and low-cost resolution.
Scalability	Low. Each transaction requires significant manual legal and administrative overhead.	High scalability for token transfers, but lacks scalable mechanisms for quality control or rights enforcement.	High. The use of ZKML for automated quality verification and smart contracts for escrow allows the system to scale efficiently.

2) *Node Economic Model and LAT Requirements:* To ensure a sustainable and reliable network of service providers, DIP Nodes are compensated through a clear economic model. Nodes register their capabilities and pricing structures on-chain, allowing for a transparent service marketplace. The primary revenue streams are:

- **Service Fees:** Direct payments for computational work, such as processing a decryption request or generating a proof.
- **Data Hosting Fees:** Long-term revenue, paid by Expert Guilds from their collected Data Retention Fees, to ensure the ongoing, persistent, and secure storage of data assets. This provides a stable, recurring revenue stream that incentivizes node reliability and longevity.

For certain high-value use cases, particularly within the RFD framework, an additional economic requirement may be imposed on the service-providing node. A Buyer or Guild issuing an RFD can specify that the chosen DIP Node must itself hold the relevant LATs to be eligible to service the request. This creates a powerful, skin-in-the-game dynamic: to be entrusted with managing a specific class of high-value data, a node must first be an economic stakeholder in that same data. This requirement can be programmatically enforced by the RFD smart contract.

B. Security and Consensus: An L2-Native Approach

The DIP Protocol is designed to be architecturally agnostic and deployable on any sufficiently advanced EVM-compatible Layer 2 (L2) network. It inherits the security and consensus mechanism of its host chain. The network's core security is

not maintained by the DIP Nodes themselves, but rather by the standard validators of the L2, who stake the host chain's native asset (e.g., ETH).

This deliberately decouples the role of network security from the role of specialized service provision. DIP Nodes are economic actors who compete in a free market to provide privacy and data management services, while network validators are responsible for the integrity of the underlying ledger.

C. A Decentralized Key Management Scheme via Multi-Party Computation (MPC)

A critical security challenge for any system that manages encrypted data is the management of decryption keys. A centralized key holder would represent a single point of failure and a prime target for attack. The DIP Protocol addresses this challenge through a sophisticated, decentralized key management architecture that utilizes Multi-Party Computation (MPC) [34]. MPC is a cryptographic subfield that allows multiple parties to jointly compute a function over their inputs while keeping those inputs private. In DIP, it is used to manage decryption keys without any single party ever holding a complete key [35]–[38].

The protocol employs a two-key MPC scheme to balance efficiency with security:

- **Key Share 1 (Node-Specific Key):** For any given RFD, the primary decryption key is secret-shared. A crucial share is held by the specific DIP Node selected by the Buyer. This allows for highly efficient and low-cost data access under normal operating conditions. The node operator is held accountable through a slashing mechanism; if they can be proven to have colluded in

disclosing data to an unauthorized party, their stake is forfeited.

- **Key Share 2 (Network Consensus Key):** A second, complementary share of the decryption key is distributed among the broader consensus of DIP network participants (or a designated committee thereof). This network key is not used for routine data access. It serves as a high-reliability, decentralized backup mechanism. In the event that the primary DIP Node becomes unavailable or maliciously withholds data, the LAT holder can petition the network. If the request is validated by consensus, the network can use its key share to reconstruct the full key and grant the holder access to their data.

This dual-key architecture provides the best of both worlds: the efficiency of a single, designated operator for normal operations, and the robust security and liveness guarantees of a fully decentralized consensus mechanism for fault tolerance and recovery [39], [40].

D. Solving the Disclosure Problem: A Verifiable Quality Function Interface

The technical core of the DIP Protocol, and its solution to the data disclosure problem, lies in establishing a standardized interface for a **verifiable quality function**. This is a smart contract interface that allows a Buyer to programmatically define the conditions a dataset must meet. The protocol remains agnostic to the specific verification method, allowing for a spectrum of trust models to coexist in a competitive marketplace. The interface is designed to support two primary modes of verification:

- **Pre-Purchase (Trustless) Verification:** This model solves the data disclosure problem completely by allowing for verification **before** the data is revealed to the Buyer. This is the most secure method, ideal for high-value IP exchange in fully trustless environments. The primary implementation for this is Zero-Knowledge Machine Learning (ZKML), as detailed below.
- **Post-Purchase (Dispute-Based) Verification:** For lower-stakes transactions or in cases where the Buyer and Seller have a pre-existing degree of trust, a simpler model can be used. The data is delivered to the Buyer upon payment, but the funds are held in escrow for a defined period. The Buyer can then run their own quality checks. If the data is found to be deficient, the Buyer can submit proof of this deficiency to the escrow contract to initiate a dispute, which is then arbitrated by the Expert Guild.

This flexibility allows participants to choose the appropriate trade-off between cryptographic security and transactional simplicity for their specific needs.

1) Advanced Implementation: Zero-Knowledge Machine Learning (ZKML):

E. Native Opcodes for Efficient Privacy Services

If the DIP protocol is implemented as a sovereign Layer 1 or Layer 2, it can introduce novel, precompiled contracts or native EVM opcodes to dramatically improve the efficiency

and developer experience of interacting with its privacy infrastructure. These opcodes would serve as a standardized, gas-efficient interface for smart contracts to request services from the network of DIP Nodes.

For example, the protocol could include:

- **DIP_ENCRYPT_POLICY:** A precompile that allows a smart contract to submit a data payload, specify an access policy (e.g., requiring ownership of a specific LAT), and request its encryption and storage by a designated DIP Node. The precompile would handle the routing of the request to the node and the on-chain registration of the access policy.
- **DIP_DECRYPT_REQUEST:** A precompile that allows a smart contract, on behalf of a user, to verify the necessary on-chain conditions (e.g., burning a LAT) and then formally request the decryption of the corresponding data from the responsible DIP Node.

By embedding these functions at the protocol level, the cost and complexity of calling these services are significantly reduced. It abstracts away the need for developers to manage complex off-chain communication with nodes, making the integration of advanced privacy features as simple as a standard smart contract call. This deep integration makes the DIP network a uniquely powerful environment for building privacy-preserving applications.

1) Advanced Implementation: Zero-Knowledge Machine Learning (ZKML): For the highest level of assurance, the protocol's verifiable quality interface is best implemented using Zero-Knowledge Machine Learning (ZKML). This technology enables the verification of data properties without revealing the underlying data itself, making trustless commercial exchange of IP possible.

A Zero-Knowledge Proof (ZKP) is a cryptographic protocol in which one party (the Prover) can prove to another party (the Verifier) that a given statement is true, without conveying any information apart from the fact that the statement is indeed true [41]–[43]. ZKPs are characterized by three properties:

- **Completeness:** If the statement is true, an honest Prover can always convince an honest Verifier.
- **Soundness:** If the statement is false, a dishonest Prover cannot convince an honest Verifier (except with a negligible probability).
- **Zero-Knowledge:** The Verifier learns nothing other than the validity of the statement.

ZKML extends this concept to the domain of machine learning [44]–[47]. It allows a party to prove that they have correctly executed a specific ML model on a specific input to produce a specific output, all without revealing the input, the model's internal parameters (weights), or the intermediate computational steps. This is achieved by representing the ML model's operations as an arithmetic circuit and using a ZKP system, such as a zk-SNARK (Zero-Knowledge Succinct Non-Interactive Argument of Knowledge), to generate a compact, efficiently verifiable proof of the computation's integrity [48]–[50].

Within the DIP Protocol, ZKML is the mechanism for objective and automated quality assurance. In an RFD, a Buyer can specify not just a subjective description of the data they want, but a precise, computational "quality function." This function can be an algorithm, a set of logical rules, or, most powerfully, an ML model designed to detect specific properties (e.g., a model that classifies images for content, a model that checks text for toxic language, or a statistical model that verifies a dataset's distribution and lack of bias).

A data contributor (Worker) can then take their private dataset, run it through the Buyer's public quality model, and generate a ZK proof. This proof cryptographically attests to the statement: "When my private dataset is used as input to the specified public quality model, it produces the claimed output (e.g., 'passes quality check' with a score of 99%)." The Worker can submit this proof to the escrow smart contract. The contract, and the Buyer, can efficiently verify the proof on-chain without ever gaining access to the Worker's proprietary dataset [51]–[53]. This process makes the arbitration role of the Expert Guild objective and scalable. Instead of mediating subjective disputes, the Guild's primary function becomes the verification of cryptographic proofs, transforming a complex social problem into a straightforward computational one. This verifiable computation model replaces reputation-based trust with mathematical certainty, a paradigm shift that is essential for building a truly trustless data economy [54]–[56].

F. The Request-for-Data (RFD) Lifecycle: A Technical Walkthrough

The integration of these components can be illustrated by a step-by-step walkthrough of a complete transaction on the DIP Protocol:

- 1) **RFD Issuance:** An AI company (the Buyer) wishes to acquire a dataset of 10,000 expertly annotated medical images for training a diagnostic model. The Buyer creates an RFD smart contract with the following parameters:
 - Data Specification: "10,000 chest X-ray images, annotated for pneumonia."
 - Quality Level: Requires LATs of quality level $q \geq 0.9$.
 - Trusted Node: Specifies the address of a reputable DIP Node for data hosting and decryption.
 - ZKML Verifier: Specifies the address of a smart contract containing a pre-trained ML model that verifies the accuracy of annotations and flags images with artifacts.
 - Escrow: The Buyer deposits payment (e.g., in USDC) into the RFD contract.
- 2) **Data Generation and Proof Creation:** A community of radiologists (the Experts) collaborates via a CVPP instance to annotate a large pool of X-ray images. Through this meritocratic process, they earn LATs representing their contributions at various quality levels. An expert who has earned a sufficient number of LATs at $q \geq 0.9$

compiles a dataset of 10,000 images. They then use the Buyer's specified ZKML verifier model to generate a zk-SNARK proof, demonstrating that their dataset passes the quality check.

- 3) **Escrow Fulfillment:** The Expert submits their LATs and the generated ZK proof to the Buyer's RFD contract.
- 4) **Verification and Access Grant:** The RFD smart contract executes the verifier function on-chain, which confirms the validity of the ZK proof. This action is computationally inexpensive. Upon successful verification, the contract transfers ownership of the LATs to the Buyer.
- 5) **Decryption and Data Delivery:** The Buyer presents their ownership of the LATs to the trusted DIP Node specified in the RFD. The node verifies the on-chain ownership record and, using its MPC key share, decrypts the underlying image dataset and delivers it securely to the Buyer.
- 6) **Settlement:** The RFD contract, having confirmed the successful verification and token transfer, automatically releases the USDC payment to the Expert. Per the contract's terms, the transferred LATs are then sent to a burn address, permanently removing them from circulation and eliminating any risk of governance dilution for the Expert Guild.

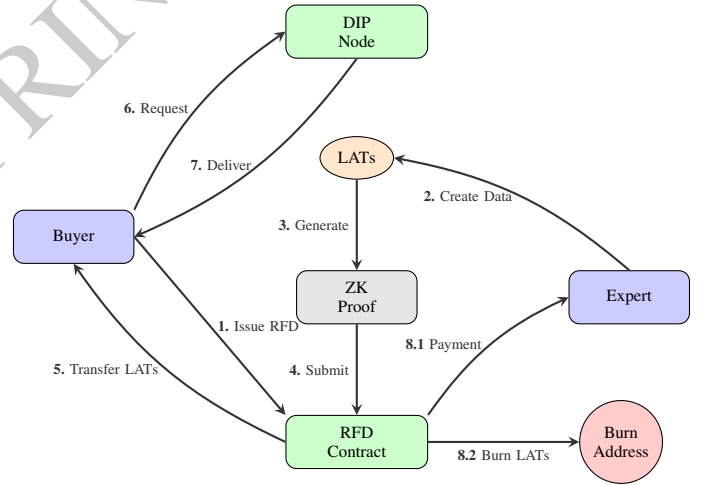


Fig. 2. The Request-for-Data (RFD) Lifecycle within the DIP Protocol, from issuance to final settlement.

V. APPLICATIONS AND ECOSYSTEM

A. Powering the Next Generation of AI Models with Verifiable Data

The primary and most immediate application of the DIP Protocol is to serve the burgeoning AI and ML industry. AI development companies can leverage DIP as a decentralized marketplace to source bespoke, high-quality datasets with unprecedented guarantees. Through the RFD process, a company can commission the creation of training data that is cryptographically verified to meet specific criteria, such as:

- **Bias Mitigation:** Using a ZKML model to prove that a dataset has a balanced demographic distribution, helping to create fairer and more equitable AI systems.
- **Content Filtering:** Ensuring a dataset is free of harmful, toxic, or copyrighted content by requiring it to pass a ZKML classifier model.
- **Provenance Verification:** Guaranteeing that a dataset is composed entirely of verifiably human-generated data, providing a crucial defense against model collapse and subliminal learning from synthetic data [51].

This moves the industry beyond relying on the reputation of data vendors to a new standard of mathematically verifiable data integrity, fundamentally de-risking the AI development lifecycle.

B. Creating Sustainable Markets for Knowledge Work

The framework established by the DIP Protocol extends far beyond datasets. It can be applied to any form of digital knowledge work where quality can be objectively or semi-objectively assessed. This transforms the "Knowledge work as a service" concept from a theoretical model into a practical reality [7]. Potential applications include:

- **Decentralized Security Audits:** A DAO could issue an RFD for a security audit of its smart contracts, specifying a ZKML verifier that runs a suite of automated analysis tools. Auditing firms could compete to find vulnerabilities and submit their findings, with payment escrowed and released upon successful verification.
- **Verifiable Academic Research:** A research foundation could fund studies by requiring researchers to submit their results along with ZK proofs demonstrating that their statistical analysis was performed correctly on a private dataset, enhancing reproducibility and trust in scientific findings.
- **Creative Content Generation:** A gaming company could commission the creation of in-game assets (e.g., character designs, lore) through a CVPP process, with the community of players and artists voting on the best submissions, ensuring the final content is aligned with player preferences.

C. The Meritocratic Flywheel: Unifying Digital Marketplaces

The principles underpinning the DIP Protocol can be generalized beyond the data economy to create a powerful new model for any digital marketplace, from ride-sharing and food delivery to freelance services. This is achieved by abstracting ACIP's "energy input" from a direct financial stake to any verifiable, value-additive contribution to the network.

Consider a ride-sharing service built as an MAO on the DIP framework. The core "energy input" (x) required to earn governance rights is no longer just a stake, but the platform's primary revenue metric: the fare. Both the driver who provides the service and the rider who pays the fare are making verifiable contributions. The protocol can be configured to mint governance tokens proportional to this revenue generated, allocating them to both the service provider and the consumer.

This creates a powerful economic flywheel:

- 1) **Ownership attracts providers:** Drivers are strongly incentivized to join the platform that grants them a share of ownership and a say in its governance, a benefit no traditional, centralized competitor can offer.
- 2) **More providers improve service:** A larger pool of drivers leads to shorter wait times and better service, attracting more riders.
- 3) **More users create more value:** Increased platform usage generates more revenue, which in turn mints more governance rights, reinforcing the ownership stake of all active participants.

This flywheel effect, where usage directly translates into ownership, creates a network effect that is fundamentally more defensible than the purely capital-based moats of Web2 marketplaces.

D. The End of App Monopolies: A Unified User Experience

This MAO-based model has a profound implication for the end-user: the unification of competing services into a single, seamless user experience.

Even if multiple competing teams or companies build different front-end applications for the same service (e.g., three different ride-sharing apps), they would all operate on the shared DIP back-end. Service providers (the drivers) would naturally gravitate to the application that offers them the best terms and the most governance, but their services would be available across the entire network.

This leads to a paradigm shift:

- **For Users:** Instead of needing multiple apps for the same service, a user can access all providers through a single interface. Competition shifts from building walled-garden monopolies to providing the best user interface, discovery algorithms, or customer service on top of a shared, open protocol.
- **For Businesses:** The opportunity for entrepreneurs is preserved, but the model changes from platform monopolies to what can be described as "time-duplexed" competition. Different teams can compete for market share at different times, but they cannot lock in users or providers permanently. The underlying protocol and its pool of providers and users remain a shared public good.

This creates a market that is simultaneously more competitive for businesses and simpler and more efficient for consumers, aligning incentives to prevent the extractive monopolies that characterize the current digital landscape.

E. Enhancing DAO Governance and R&D through Meritocratic Tooling

The core primitives of the DIP Protocol, particularly ACIP and CVPP, are powerful governance tools in their own right and can be adopted by other DAOs to address common challenges. Many DAOs suffer from plutocratic governance models (1-token-1-vote) and low participation rates. By integrating ACIP, a DAO can identify and empower members

with demonstrated expertise and contribution, moving beyond simple token-based voting towards a more meritocratic and effective governance system [7], [57], [58]. For instance, a protocol DAO could use ACIP to create a "Core Developer" guild, where only members who have achieved a certain competence rank in software development can vote on technical upgrade proposals. This ensures that critical decisions are made by the most qualified participants, rather than just the wealthiest.

F. On-Chain Compliance and Verifiable Decentralization

A critical, often unacknowledged, weakness of major blockchain ecosystems is the illusion of decentralization. While protocol rules are decentralized, the physical node infrastructure is often highly centralized. A majority of nodes for networks like Ethereum and Solana are hosted in a small number of data centers within a single jurisdiction (primarily the United States). This geographic concentration creates a systemic risk, making the network vulnerable to nation-state level regulation, censorship, or physical disruption, and severely limits its appeal for global institutional use cases that require jurisdictional guarantees.

The DIP Protocol is designed to solve this problem by treating decentralization not just as a software property, but as a verifiable, physical, and jurisdictional attribute. As described in Section 4.1, the protocol's service-oriented model inherently incentivizes the creation of a globally distributed network of nodes that compete on security, performance, and jurisdictional compliance.

Building on this foundation, the protocol can be implemented as a sovereign L1 or L2 that provides a native solution for on-chain compliance. This is achieved through an additional, optional field in the protocol's transaction structure designed for **aggregated compliance attestations**. This field allows a transaction to be co-signed by a specific quorum of registered DIP Nodes.

This seemingly simple feature unlocks a powerful new paradigm for regulated institutions:

- **Built-in Compliance Primitives:** Monetary authorities and financial institutions no longer need to build private blockchains. They can simply run a registered and verified DIP Node (or a quorum of nodes) within their jurisdiction. Their signature in the attestation field becomes a verifiable, on-chain "stamp of approval" for any transaction.
- **Enforceable Smart Assets:** A smart wallet or a specific token contract (e.g., a regulated stablecoin) can be programmed to require that all its transactions include a valid signature from a designated jurisdictional node. This allows for the creation of assets that are, by design, compliant with specific regulatory regimes, without compromising the public and permissionless nature of the broader network.

This architecture provides a direct path for bridging the gap between the decentralized world and institutional finance. It moves beyond the false dichotomy of permissionless versus permissioned systems, offering a network that is publicly

accessible but allows for the creation of verifiable, opt-in "compliant enclaves" for specific assets and applications.

G. A Semi-Open Source Paradigm for Digital Public Goods

The DIP Protocol enables a novel and sustainable economic model for digital public goods, formalizing the "semi-open source" paradigm [7]. The current Web3 landscape is dominated by a binary choice: either a project is fully open-source and struggles with monetization, or it is fully proprietary and centralized. DIP creates a middle ground.

Using LATs, the creators of a valuable piece of IP (e.g., a dataset, a software library, a research paper) can implement tiered and programmable access. Commercial entities wishing to use the IP for profit can be required to purchase LATs on the open market, providing a direct and sustainable revenue stream to the creators. Simultaneously, the creators can grant free or heavily discounted access to non-commercial users, such as academics, students, or hobbyists. This model can be enforced through smart contracts, creating a fair and transparent system where those who derive commercial value from public goods contribute to their upkeep. This provides a robust solution to the chronic underfunding of open-source infrastructure and contrasts sharply with the often ambiguous and legally complex licensing models currently used for NFTs and other Web3 assets [59]–[62].

VI. CONCLUSION: TOWARDS A TRUE OWNERSHIP ECONOMY FOR INTELLECTUAL PROPERTY

The digital revolution, while unlocking unprecedented access to information, has simultaneously devalued it by making perfect, costless replication trivial. This has led to a fundamental market failure in the valuation and protection of intellectual property, a problem that has become acutely urgent with the rise of a data-driven AI economy. The Decentralised Intellectual Property (DIP) Protocol offers a comprehensive, first-principles solution to this long-standing challenge.

By integrating a meritocratic system for data provenance (ACIP/CVPP), a game-theoretically sound economic architecture (LATs and Expert Guilds), and a privacy-preserving technical infrastructure (MPC/ZKML), DIP re-establishes the concept of verifiable scarcity for digital assets. It moves beyond the limitations of traditional legal frameworks and the nascent, often inadequate, IP models of Web3. The protocol's ability to solve the data disclosure problem—allowing for the verification of an asset's properties without revealing the asset itself—is the critical cryptographic breakthrough that unlocks a truly trustless and liquid market for IP.

DIP is more than a technical protocol; it is a foundational economic and social layer designed for a future where intellectual contributions are the primary drivers of value. It provides the tools to ensure that the creators of this value—the domain experts, the researchers, the artists, the data curators—are fairly compensated and empowered. By transforming intellectual property from an abstract legal concept into a concrete, liquid, and sovereign digital asset, the DIP Protocol lays the

groundwork for a more equitable, efficient, and innovative digital ownership economy.

REFERENCES

- [1] Helios Solutions, *The data foundation: Why data quality is crucial for ai/ml success*, n.d. [Online]. Available: <https://www.heliossolutions.co/blog/the-data-foundation-why-data-quality-is-crucial-for-ai-ml-success/>.
- [2] AIMultiple, *Data quality in ai: Challenges, importance & best practices*, n.d. [Online]. Available: <https://research.aimultiple.com/data-quality-ai/>.
- [3] Keymakr, *Future trends in data quality: Ai and machine learning*, Sep. 2024. [Online]. Available: <https://keymakr.com/future-trends-in-data-quality-ai-and-machine-learning/>.
- [4] Stanford HAI, *The 2025 ai index report*, 2025. [Online]. Available: <https://aiindex.stanford.edu/>.
- [5] McKinsey, *The state of ai: How organizations are rewiring to capture value*, n.d. [Online]. Available: <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2023-generative-ai-breakout-year>.
- [6] Northwest Executive Education, *Increasing demand for artificial intelligence – 5 key factors*, May 2025.
- [7] T. Pecherskis, A. Smirnovs, and A. Soboleva, *Continuous voting-proposing protocol for ordering group intents*, 2024.
- [8] The Times of India, *Ai models may secretly pass on hidden behaviours, warns study*, Jul. 2025.
- [9] HackerNoon, *To hit \$1t tvl, ethereum must play the ace*, n.d. [Online]. Available: <https://hackernoon.com/to-hit-dollar1t-tvl-ethereum-must-play-the-ace-that-is-data-sovereignty-194z37a5>.
- [10] TokenMinds, *5 steps to protect intellectual property in web3 world*, Feb. 2024. [Online]. Available: <https://tokenminds.co/blog/protect-intellectual-property-in-web3/>.
- [11] Endless Domains, *Web3 and legal frameworks: Navigating intellectual property, smart contracts, and regulatory challenges*, Feb. 2025.
- [12] Asia IP Law, *How web 3.0 revolutionizes intellectual property: Nfts, daos and decentralized ownership*, Jan. 2024.
- [13] Zama, *Announcing the zama confidential blockchain protocol*, n.d. [Online]. Available: <https://www.zama.ai/post/announcing-the-zama-confidential-blockchain-protocol>.
- [14] Wikipedia, *Full disclosure (computer security)*, n.d. [Online]. Available: [https://en.wikipedia.org/wiki/Full_disclosure_\(computer_security\)](https://en.wikipedia.org/wiki/Full_disclosure_(computer_security)).
- [15] Wallarm, *What is an information disclosure? examples and prevention*, n.d. [Online]. Available: <https://www.wallarm.com/what/information-disclosure>.
- [16] National Center for Biotechnology Information, *Preserving privacy using technology from computer science, statistical methods, and administrative procedures*, n.d.
- [17] American Statistical Association, *D: Computer science approaches*, n.d.
- [18] National Bureau of Economic Research, *Valuing the us data economy using machine learning and online job postings*, n.d.
- [19] The Times of India, *Data science vs. economics: Which degree is better suited for a data-driven future?* Jul. 2025.
- [20] SG Analytics, *New data economy – future of data-driven value creation*, May 2025.
- [21] Astera, *What is data provenance?* n.d.
- [22] Identity.com, *What is data provenance?* n.d.
- [23] T. Pechersky and A. Smirnovs, *Autonomous competence identification protocol: A dynamic ranking ladder system for blockchain applications*, 2025.
- [24] Peeramid Network, *Rankify protocol: A framework for decentralized intellectual property management*, n.d.
- [25] *Continuous voting proposing protocol (cvpp) description*, n.d.
- [26] Rapid Innovation, *Tokenomics guide: Mastering blockchain token economics 2024*, 2024.
- [27] TokenMetrics, *What is tokenomics? a beginner's guide*, n.d.
- [28] Rock'n'Block, *Tokenomics decoded: Key components explained*, n.d.
- [29] CiteSeerX, *A game-theoretic model of third-party agents for enforcing obligations in transactions*, n.d.
- [30] H. Zhang, *Hope or hype: On the viability of escrow services as trusted third parties in online auction environments*, n.d.
- [31] ResearchGate, *Barriers to collusion-resistant transaction fee mechanisms*, n.d.
- [32] Frontiers in Blockchain, *Collusion resistance of voting systems in decentralized autonomous organizations*, 2024.
- [33] Forbes Business Development Council, *Game theory in blockchain: The invisible hand of decentralized trust*, Jul. 2025.
- [34] Wikipedia, *Secure multi-party computation*, n.d.
- [35] Cyfrin, *Multi-party computation: Secure private collaboration*, n.d.
- [36] GeeksforGeeks, *What is secure multiparty computation?* n.d.
- [37] Decipher, *Secure multi-party computation*, n.d.
- [38] Chainlink, *Secure multi-party computation*, Aug. 2024.
- [39] Cointelegraph, *Mpc wallets explained*, Aug. 2024.
- [40] Wipro, *Synergizing blockchain and multi-party computation to reimagine transactions*, n.d.
- [41] Circularise, *Zero-knowledge proofs explained in 3 examples*, n.d.
- [42] Wikipedia, *Zero-knowledge proof*, n.d.

- [43] Identity.com, *What are zero-knowledge proofs (zkp)?* Jul. 2025.
- [44] OSL, *How to verify on-chain machine learning algorithms using zero-knowledge*, n.d.
- [45] Meegle, *Zero-knowledge proof in machine learning*, n.d.
- [46] Cloud Security Alliance, *Leveraging zero-knowledge proofs in machine learning and llms: Enhancing privacy and security*, Sep. 2024.
- [47] A. Sharma, *Zero-knowledge machine learning: Bridging privacy and verification in ai systems*, n.d.
- [48] A. Nitulescu, *A gentle introduction to zk-snarks*, n.d.
- [49] D. Kang, *Trustless machine learning inference*, Oct. 2022.
- [50] S. Thakur and J. Breslin, *Efficient deep neural network verification with qap-based zksnark*, n.d.
- [51] Datawallet, *What is zkml?* n.d.
- [52] Vara Network, *Zkml: Empowering privacy on vara network*, n.d.
- [53] Ledger, *Zero-knowledge machine learning*, Mar. 2025.
- [54] Archetype, *Verifiable compute: Scaling trust with cryptography*, n.d.
- [55] ResearchGate, *Verifiable computing applications in blockchain*, 2021.
- [56] W. Zhang, *Verifiable ai on bitcoin*, n.d.
- [57] Colony, *What is reputation-based governance?* n.d.
- [58] Rapid Innovation, *Dao governance models explained: Token-based vs. reputation-based systems*, n.d.
- [59] Solana Compass, *Ip licensing*, n.d.
- [60] Neptune Mutual, *Protecting your intellectual property in web3*, n.d.
- [61] BlockApps, *Protecting intellectual property in web3*, n.d.
- [62] ICON Partners, *Strategies for web3 intellectual property management*, n.d.