Origin: The IP Operating System for the Autonomous Economy

Transforming IP Ownership for the AI-Native Era

The Origin Whitepaper

Camp Foundation

Abstract

The rise of autonomous AI agents has revealed a structural failure in traditional intellectual property systems: they cannot operate at machine speed, handle algorithmic-scale reuse, or enforce rights in environments where content is created and transformed continuously. Origin is a programmable intellectual property protocol that makes rights machine-readable and automatically enforceable, enabling instant licensing, transparent attribution, and real-time royalty settlement for both human creators and AI agents.

Origin represents all IP as composable onchain assets called IPNFTs (Intellectual Property Non-Fungible Tokens), each containing executable licensing terms, derivation logic, and royalty rules. When an asset is used by an application, a human, or an autonomous agent, the protocol validates rights, executes the appropriate license, and distributes payments atomically. Powered by a high-throughput blockchain layer, Origin settles micro-royalties at sub-cent granularity while preserving cryptographic proof of ownership and usage.

By treating AI agents as first-class economic participants, Origin provides the foundational infrastructure for an AI-native economy where the creation, composition, licensing, and monetization of intellectual property occur autonomously and at global scale.

Contents

1 Introduction

The emergence of autonomous AI agents as content creators represents a phase transition in intellectual property management. These systems now generate millions of creative works daily, remixing existing content, collaborating with other agents, and building derivative works at speeds that make human oversight impossible. A single AI agent might reference hundreds of source materials in a generated image, compose music from dozens of training samples or synthesize text from thousands of documents at latencies measured in milliseconds, far beyond any process that relies on human verification, contractual negotiation, or manual rights clearance. Traditional IP systems, built on manual registration and negotiated licensing, cannot track attribution or enforce rights at this scale or speed.

Current blockchain-based solutions fail to address these requirements. Existing NFT standards like ERC-721 and ERC-1155 represent ownership but lack the semantic structure to encode licensing terms, attribution requirements, or royalty distribution logic. Layer 1 networks with transaction costs above $0.10 make micro-royalty payments economically infeasible. As an example, when a derivative work owes $0.001 to each of 100 source creators, settlement costs exceed the payments by two orders of magnitude. Most critically, these systems provide no mechanism for legal enforceability beyond proof of registration, meaning they operate as passive catalogs rather than active rights engines capable of enforcing usage, validating permissions, or settling obligations.

To support AI-scale creativity, IP infrastructure must satisfy four requirements: (1) deterministic attribution, (2) programmable licensing, (3) sub-cent economic settlement, and (4) real-world legal enforceability. Origin is engineered specifically to meet these constraints.

Origin addresses these limitations through three core architectural innovations:

1. **Programmable Rights**: IPNFTs extend the NFT primitive with embedded licensing logic and royalty distribution rules, making IP rights machine-readable and automatically executable.

2. **Attribution Graph**: A provenance graph tracks derivation relationships between assets, enabling royalties to flow automatically from derivative works back through chains of attribution to original creators.

3. **Scalable Economics**: Integration with Camp Network's high-throughput infrastructure enables batch settlement of micro-royalties at costs well below $0.001 per transaction, making it economically viable to compensate creators for even fractional usage of their work. Together, these components create a programmable IP layer where AI agents and humans can create, license, and monetize intellectual property autonomously and economically.

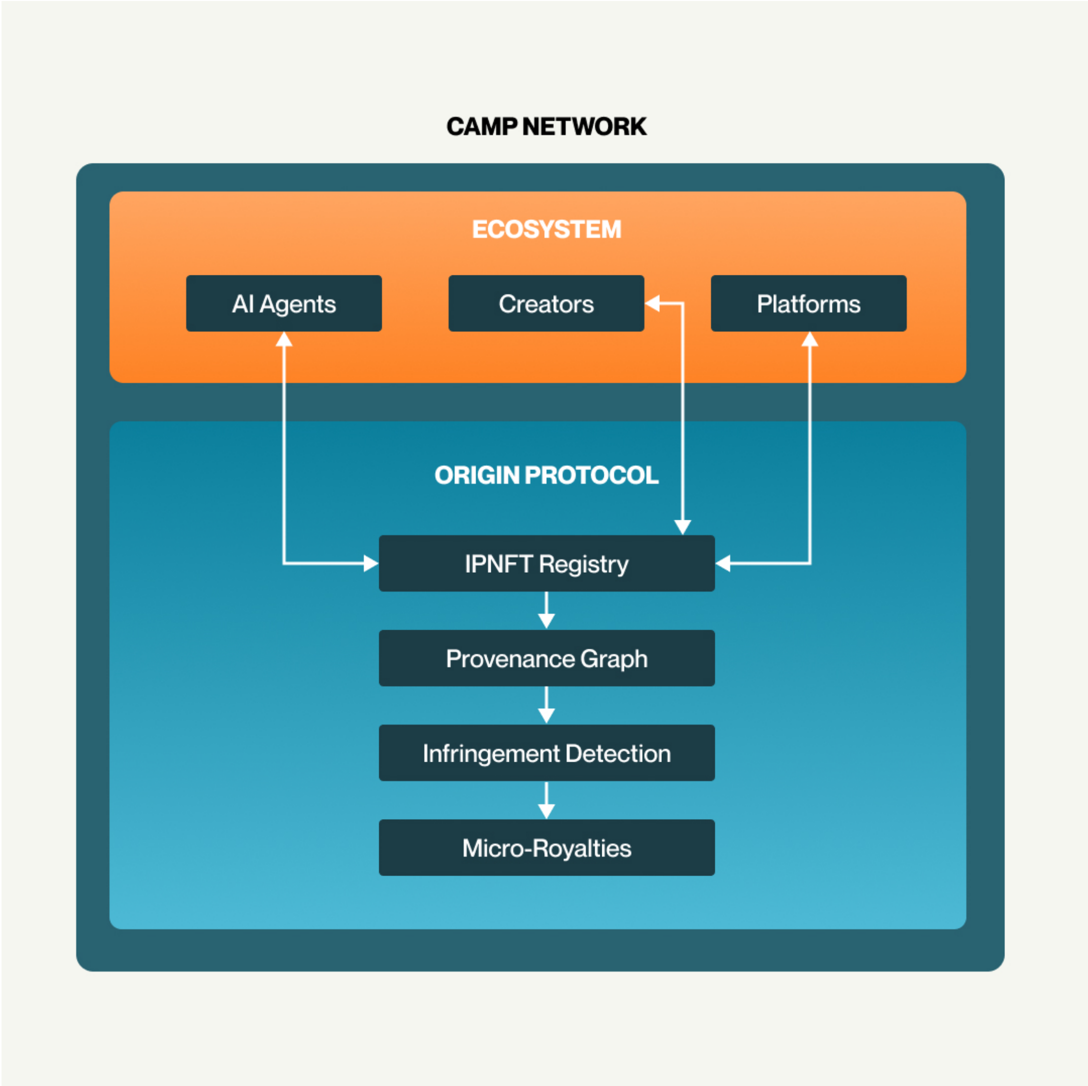

Figure 1: Origin Ecosystem Interaction Map

These three primitives together form a vertically integrated operating system for intellectual property, enabling autonomous creation, autonomous licensing, and autonomous monetization.

At AI scale, conflicts are inevitable. Multiple parties may claim authorship, agents may misuse IP under the wrong license, or derivatives may exceed permitted similarity thresholds. To solve this, Origin includes an onchain Dispute Module that permits users to report suspected infringement and halt the corresponding IPNFT during the evaluation period. The protocol records relevant evidence (hashes, timestamps, signatures, similarity proofs) and routes the claim through a deterministic resolution flow. This creates an enforceable mechanism where disputes can be resolved with cryptographic evidence rather than subjective interpretation, aligning the system with both legal and machine requirements. Once resolved, Origin updates ownership, attribution, and royalty flows accordingly. This ensures that programmable IP remains enforceable even when creation happens autonomously and at machine speed.

The fundamental opacity of modern AI systems has created a provenance crisis that threatens creative attribution. Current Transformer models from Claude to Sora are trained on user data and IP without transparent opt-in systems, making attribution impossible and, as a result, preventing creators from being compensated, enforce licensing terms, or establish that infringement occurred. As model complexity increases, tracing influence becomes computationally harder while the economic stakes grow larger, making cryptographic provenance a prerequisite for sustainable AI ecosystems. This architectural opacity represents a systemic failure to maintain the chain of custody that underpins intellectual property rights.

Scaling compounds the problem. When thousands of AI agents generate content every minute, manual infringement detection becomes structurally impossible - copyright holders cannot review every generated image, text, or audio file. Traditional takedown systems like DMCA operate on complaint-based models that assume human-scale creation rates and manual review processes. These mechanisms are inherently non-scalable relative to current and future AI content generation speeds.

Origin's provenance graph makes infringement detection cryptographically verifiable and computationally cheap. Every IPNFT carries an onchain record of its licensing status and derivation history. When an AI agent creates derivative work, the content is hashed and matched against registered IP for similarity detection. Upon identifying potential source materials, the system checks for valid licensing transactions - if none exist, the content is automatically flagged for infringement. Authorized derivatives register their source IPNFTs onchain, creating an immutable attribution chain.

This architecture shifts enforcement from reactive, human-driven complaint handling to proactive, automated verification executed at the speed of computation. Marketplaces, platforms, and AI systems query the provenance graph in real-time to verify licensing compliance before distributing or monetizing content, enabling automated enforcement at the speed and scale of AI content generation.

2 System Architecture

2.1 Overview

Origin's architecture enables autonomous IP management through three interconnected layers that handle the complete lifecycle of intellectual property from creation and registration through licensing and royalty distribution.

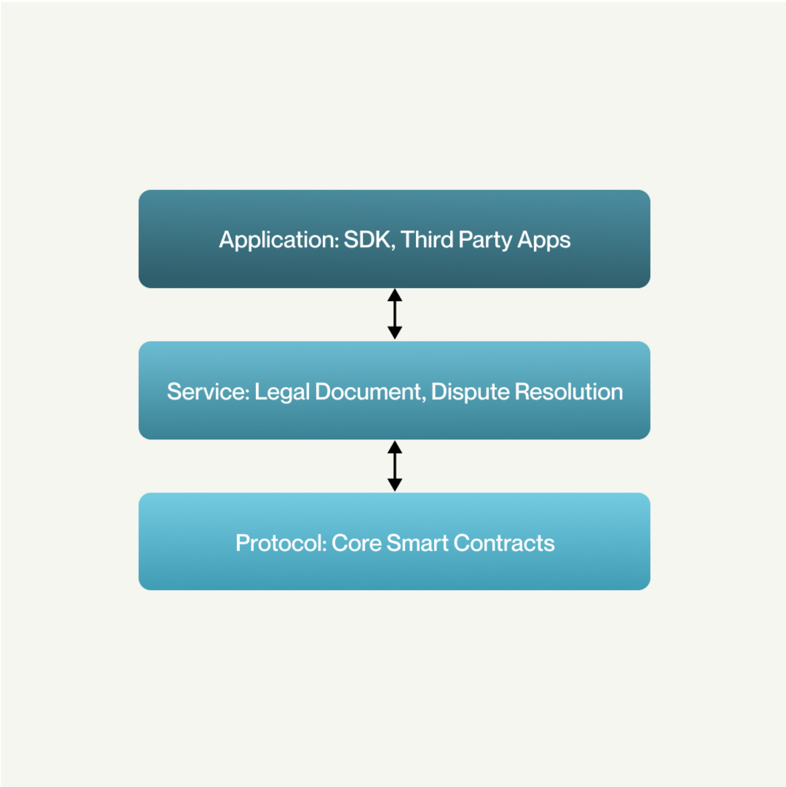

Figure 2: Origin System Overview Layers

1. **Protocol Layer**: Core smart contracts that manage IPNFT registration, license issuance, and royalty distribution. These contracts define the fundamental operations and state transitions for intellectual property assets.

2. **Service Layer**: Offchain infrastructure that provides content delivery, legal document generation, and dispute resolution services. This layer bridges the gap between onchain verification and real-world enforcement.

3. **Application Layer**: Developer tools, SDKs, and interfaces that enable third-party applications to interact with Origin's IP infrastructure. Applications range from creative marketplaces to AI training platforms.

This separation allows Origin to achieve the throughput required for AI-scale content generation while maintaining cryptographic proof of all IP relationships and transactions. This layered structure also ensures upgradeability and modularity: each layer can evolve independently as legal standards, AI capabilities, or blockchain performance constraints change without breaking the guarantees of the underlying IP model.

2.2 Design Principles

Origin's architecture follows four key principles that distinguish it from general-purpose blockchain platforms:

1. **IP-Native Design**: Rather than retrofitting token standards designed for art or collectibles (ERC-721, ERC-1155), Origin implements primitives purpose-built for intellectual property operations. IPNFTs contain embedded licensing logic, multi-party royalty splits, and derivation relationships as core properties, enabling automated IP management without external legal interpretation. This avoids forcing IP workflows into abstractions that were never designed for ownership chains, multi-party revenue splits, or legal enforceability.

2. **Composable Operations**: All protocol components expose standardized interfaces that enable complex IP workflows through simple composition. An AI agent can query the provenance graph, license multiple source assets, create a derivative work, and register the result with automatic attribution all through a sequence of composable protocol calls. These operations are intentionally deterministic, allowing agents to assemble licensing and attribution pipelines with predictable results at scale.

3. **Deterministic Execution**: Critical operations such as royalty calculation, license validation, and attribution tracking follow deterministic algorithms encoded in smart contracts. This ensures that all participants whether human creators, AI agents, or applications, observe consistent outcomes regardless of when or where operations are executed. Determinism is essential for autonomous agents: identical inputs must always yield identical rights, fees, and outcomes or automated systems cannot reason about licensing validity.

4. **Selective Decentralization**: The protocol maintains sovereignty over ownership, licensing, and provenance using on-chain smart contracts, while integrating optional off-chain services through cryptographically verified bridges. By decentralizing what must be objective (ownership, permissions, rights) and externalizing what must be subjective (legal agreements, dispute evidence, content delivery), Origin achieves both cryptographic assurance and real-world enforceability.

2.3 Architectural Requirements

To operate at AI scale, Origin's architecture is engineered around five system-level requirements:

1. **Machine-speed licensing**: Contracts must respond within milliseconds to support autonomous agent workflows
2. **Sub-cent settlement**: Royalty routing must remain economically viable for micro-derivatives
3. **Deterministic provenance**: Attribution must be verifiable without relying on heuristics
4. **Legal compatibility**: Onchain records must map cleanly into enforceable agreements
5. **Autonomous composability**: All operations must be callable programmatically by agents

These requirements shape every protocol component and distinguish Origin from generic NFT or metadata systems.

3 IPNFTs: Programmable Intellectual Property

3.1 Data Model

IPNFTs represent intellectual property as structured, machine-readable onchain objects. Each IPNFT encodes ownership, metadata, licensing rules, royalty logic, and creative lineage within a deterministic schema optimized for autonomous agent consumption. Unlike traditional NFTs, IPNFTs function as executable rights containers rather than static metadata pointers. All economic, licensing, and attribution logic needed for autonomous use is embedded directly into the object, enabling deterministic verification and monetization without intermediaries.

3.1.1 Identity

Each IPNFT is anchored to a unique, content-addressed identifier that guarantees immutability and provenance:

- *contentHash* (bytes32): Keccak-256 hash of the raw data blob, ensuring verifiable integrity
- *tokenId*: ERC-721 identifier mapping to onchain ownership
- *ownerOf(tokenId)*: Canonical authorship reference
- *baseURI*: IPFS/Arweave pointer to the asset manifest

This identity layer ensures that any consumer human or agent can verify asset authenticity prior to use. Content-addressed identity ensures that identical assets resolve to a single canonical reference, giving agents a trustless foundation for deduplication, provenance checks, and dependency resolution.

3.1.2 Metadata

IPNFT metadata includes both human-readable and machine-readable descriptors:

- Title, category, description
- Data type (e.g., dataset, model weights, media corpus)
- Media links and structured manifest files
- Update semantics for metadata vs. immutable content

Metadata is validated through content hashing and is consistently accessible through the onchain URI, enabling deterministic downstream consumption. Machine-readable metadata fields ensure that agents interpret content category, usage type, and update semantics consistently, reducing ambiguity in automated licensing decisions.

3.1.3 Licensing Terms

Licensing rules are expressed as an onchain template stored within each IPNFT:

```
struct LicenseTerms {
    uint128 price;       // Payment amount in CAMP or ERC-20
    uint32  duration;    // Access duration in seconds
    uint16  royaltyBps;  // Creator share (0-10000)
    address paymentToken;// CAMP or ERC-20 asset
}
```

Licenses define price, time-bounded access permissions, royalty parameters, and payment assets. These terms directly power the licensing workflow through:

- *buyAccess(tokenId, periods)*
- *renewAccess(tokenId, periods)*
- *hasAccess(user, tokenId)*

Standardized license parameters allow agents to compute cost, duration, and compatibility deterministically, enabling multi-step creative workflows without human validation.

3.1.4 Royalty Configuration

Royalties are executed via a dedicated routing contract (RoyaltyRouter), which performs deterministic settlement:

- Creator payment

- Protocol fee allocation
- Optional referrer share

Settlement is performed atomically at purchase time and surfaced through canonical events:

```
event RoyaltyPaid(tokenId, royaltyAmount, creator, protocolAmount);
```

Royalty flows become transparent, indexable, and automatically auditable, allowing agents and applications to reason about financial obligations in real time.

3.1.5 Attribution Graph

Each IPNFT embeds references to its parent works, forming a verifiable attribution graph that functions as a directed acyclic graph (DAG) of creative lineage. This graph captures relationships between original works and derivatives, enabling automated royalty inheritance, provenance verification, and compatibility checks across evolving assets. By encoding lineage in a structured, machine-readable form, Origin provides a native onchain representation of influence and derivation, which is essential for transparent reuse in generative or agentic systems.

3.1.6 Operational Semantics

By combining identity, metadata, licensing, royalties, and lineage, each IPNFT functions as a self-contained and upgradeable IP object. The UUPS proxy architecture ensures that contract logic can evolve without breaking the underlying storage layout or historical references. The data model is engineered for deterministic behavior: an agent can verify the asset's integrity, inspect its licensing terms, purchase the required access, and rely on predictable royalty settlement all without human intervention. Deterministic operational semantics allow agents to verify integrity, evaluate licensing terms, acquire rights, and trigger settlement without human involvement, making IPNFTs the core primitive for autonomous IP markets.

3.2 Registration Mechanism

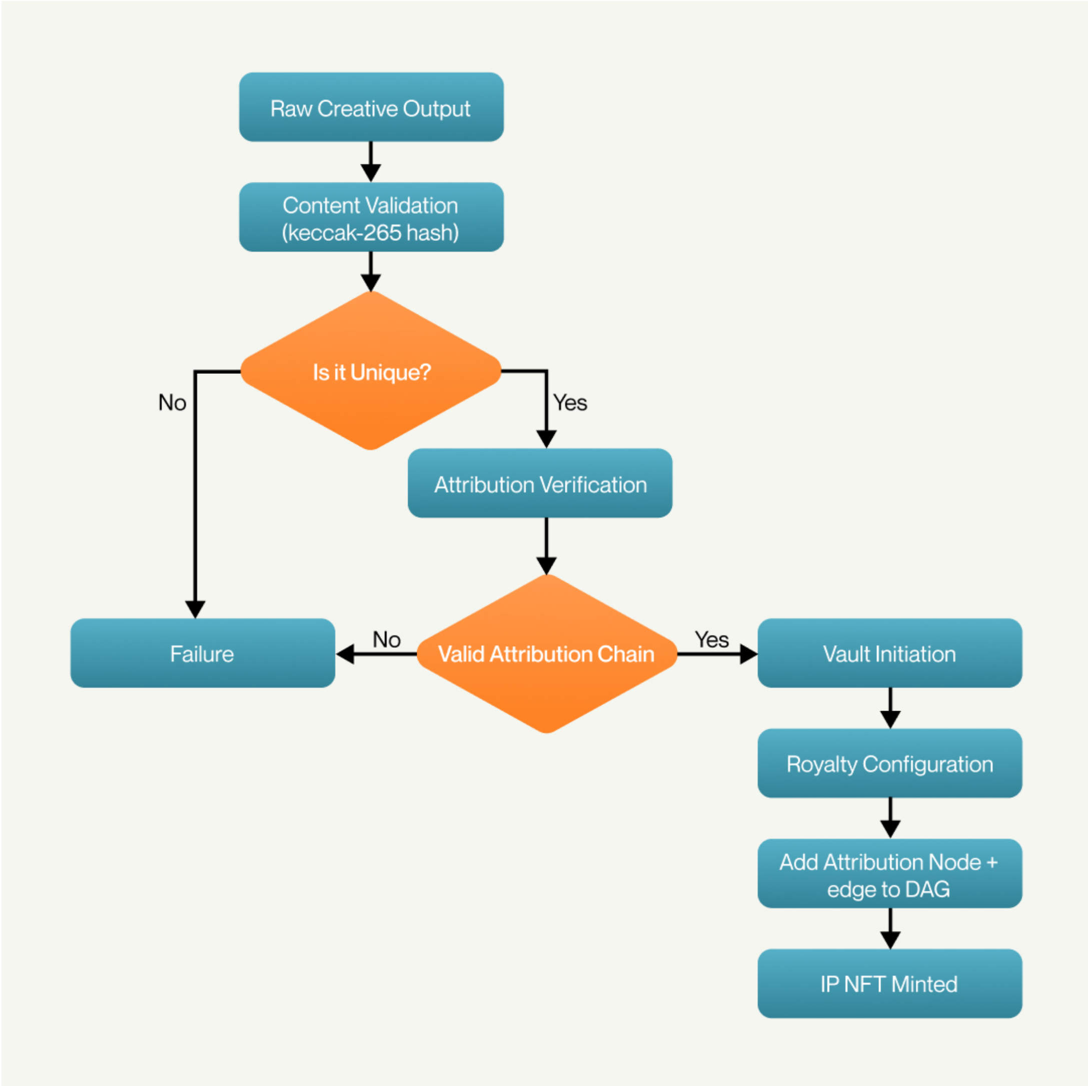

Figure 3: Origin IP Registration Flow Chart

The registration mechanism converts raw creative outputs such as datasets, media files, model weights, text corpora, and other knowledge artifacts into verifiable, programmable IP assets. Registration establishes the authoritative onchain representation of an asset and anchors it to a set of invariants that protect provenance, lineage, and economic rights.

The process begins with content validation, where the system computes a `contentHash` using a Keccak-256 digest of the submitted data and checks for prior existence. This ensures global uniqueness. Two works with identical raw content resolve to the same hash. This prevents unintentional duplication and enables deterministic referencing across applications and autonomous agents.

Next, attribution verification confirms that any declared parent works are legitimate and that the registrant has the right to create a derivative. Parent IPNFTs are referenced explicitly, allowing the protocol to validate chain of title and to enforce lineage constraints. This guarantees that derivative works inherit a correct and verifiable provenance record.

Once attribution is confirmed, the system performs vault initialization by assigning or creating a royalty routing configuration. This step links the asset to its economic infrastructure. Royalty percentages, fee splits, and settlement rules are bound to the IPNFT at creation. This ensures predictable compensation flows for both creators and upstream rightsholders.

The IPNFT is then integrated into the attribution graph, Origin's global directed acyclic graph of creative works. Each node represents an IP asset, and each edge represents a dependency or derivation. This structure supports protocol-level reasoning about provenance, license compatibility, and royalty inheritance. It also enables transparent composability across applications and agent systems.

Finally, the registration event is broadcast onchain. Indexed events allow agents, applications, and discovery services to track new assets, synchronize local graphs, and verify that the registration has been incorporated into the protocol's state.

Through this pipeline, raw data is converted into a structured and self-contained IP object with immutability, verifiable lineage, and machine-readable economic logic. This forms the foundation for automated licensing, attribution, and valuation throughout the Origin ecosystem.

3.3 Categories and Specialization

Origin supports specialized IPNFT types optimized for different asset classes:

1. **Creative Media**: Images, videos, audio with format-specific metadata
2. **Datasets**: Structured data with schema definitions and update frequencies
3. **Models**: AI models with architecture specifications and training parameters
4. **Code**: Software with dependency management and license compatibility
5. **Agents**: Autonomous entities with capability declarations and resource requirements

Each category implements domain-specific extensions while maintaining compatibility with Origin's core protocols.

4 Licensing Framework

4.1 License Taxonomy

Origin uses four primitive license scopes that applications and legal agreements can compose. Onchain, every license is the same primitive: a time-bounded access right defined by

```
LicenseTerms { price, duration, royaltyBps, paymentToken }
```

and the subscription expiry for a given wallet.

The table below separates what the protocol enforces from how that access is typically interpreted offchain.

License Scope	**Typical Permission**	**Protocol Enforcement**	**Revocation Semantics**
View	Read and display the underlying content or data	Grants time-boxed access to fetch content via `hasAccess` and the gateway. No onchain restriction on copying or redistribution once fetched.	Access ends when the subscription expires. Creator can pause new purchases or start deletion; existing subscriptions are honored until expiry.
Usage	Create transformations or derivatives that reference the IPNFT	Same as View, plus optional attribution links recorded in the attribution graph. No automatic onchain check of "fair use" or similarity, but disputes can be raised through the Dispute Module.	Same as View: time-boxed. Future access can be refused by not renewing or by entering the deletion flow; past access and settled royalties remain.
Commercial	Generate revenue from outputs that depend on the IPNFT (e.g., products, services, models)	Enforces payment at purchase time and routes royalties via the RoyaltyRouter according to royaltyBps. Off-chain agreements specify fields like territory, media, and revenue model.	Subscriptions and future commercial use rights are time-boxed. Settled payments are final; creators can change terms for future licenses or pause new access but cannot claw back past earnings.
Exclusive	Sole or priority access within a defined scope (time, territory, channel, or use case)	Expressed via onchain metadata and off-chain legal terms. Current protocol versions do not hard-block conflicting licenses; conflicts are signaled and can be challenged through the Dispute Module.	Time-boxed. Early termination or modification is handled by legal agreement or dispute outcome; at the protocol level the mechanics are identical to other licenses.

Table 1: License Taxonomy Table

4.2 Issuance Protocol

License issuance is executed through a deterministic, atomic workflow implemented in Origin's licensing and marketplace contracts. The protocol proceeds through the following stages:

1. **Authority validation**: The system verifies that the caller has the right to issue a license for the specified IPNFT, typically through creator or admin roles.

2. **Availability checks**: The protocol evaluates existing licenses to confirm that the requested license does not conflict with active or exclusive rights already in effect.

3. **Settlement**: Payment is processed using the designated token. The royalty router allocates the appropriate shares to the creator, protocol fee recipient, and any registered referrers.

4. **License creation**: A new license record is written to onchain storage with its terms, licensee address, start time, and expiry. The record includes verifiable references to the associated IPNFT.

5. **Legal binding**: If the license includes commercial or exclusive rights, a corresponding legal agreement is generated or referenced. This document relies on the onchain license state as the authoritative source of truth.

6. **State broadcasting**: Canonical events are emitted so that indexers, agents, and applications can synchronize state and track newly issued licenses.

The protocol executes atomically. Any failure in validation, settlement, or record creation triggers a full revert, ensuring that no partial or inconsistent license states can exist.

4.3 Enforcement Mechanisms

Origin applies a two-layer enforcement model that combines automated verification with legal protection. Together these layers ensure that licensing rules are respected across both blockchain-based and traditional environments.

Onchain enforcement:	**Off-chain enforcement:**
The licensing and marketplace contracts expose standardized verification functions that applications and agents can call before executing protected actions. These checks confirm that a license exists, has not expired, and is compatible with the requested operation. By relying on the license registry as the single source of truth, applications can enforce authorization deterministically and without discretionary logic.	Commercial and exclusive licenses can be accompanied by formal legal agreements generated at issuance time. These agreements incorporate the onchain license record and reference its data as the authoritative definition of rights and obligations. In the event of misuse outside the blockchain environment, the agreement provides a clear path for enforcement through conventional legal systems.

Together, these mechanisms create a unified enforcement framework that supports automatic verification for agents and decentralized applications, while also providing creators and licensees with legally enforceable protections in the off-chain world.

5 Royalty Distribution System

5.1 Attribution Graph

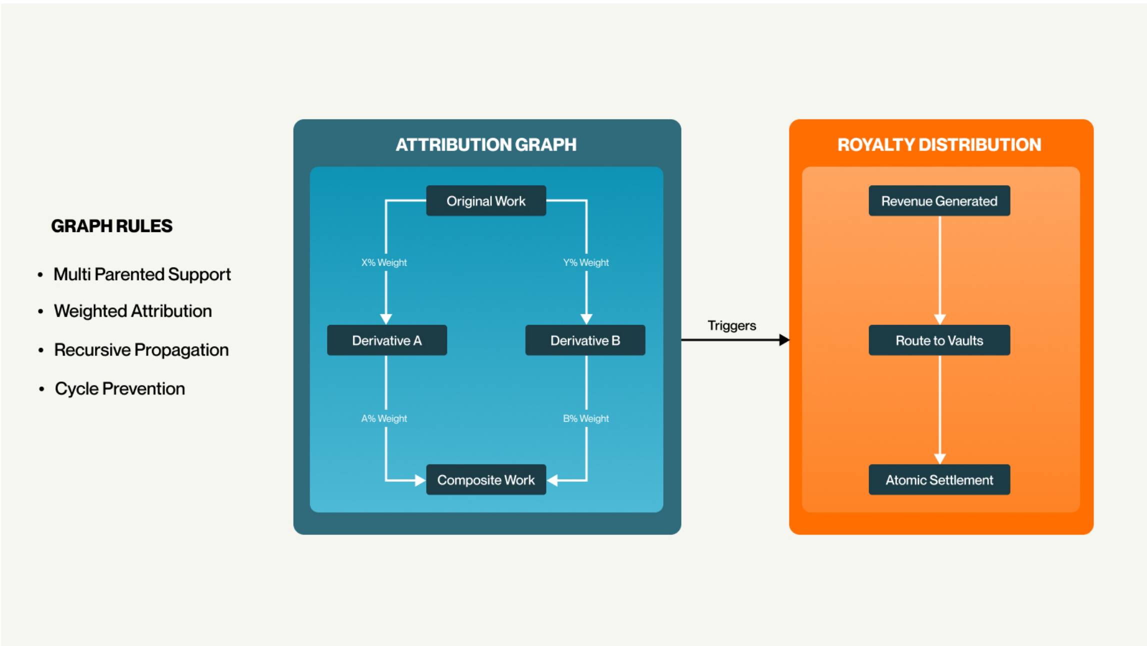

Figure 4: Camp Attribution Diagram

Origin structures provenance and revenue sharing through a parent-child attribution graph that records how each IPNFT relates to the works that influenced it. When a derivative generates revenue, the system consults this graph to determine how compensation should be allocated among upstream contributors.

The graph supports multi-parent inheritance for collaborative or composite works, weighted attribution based on declared contribution intensity, recursive propagation across lineages of any depth, and strict cycle prevention through topological ordering. This ensures that all derivations can be evaluated deterministically and that contributors receive compensation proportional to their role in the creative lineage.

5.2 Distribution Algorithm

Royalty distribution follows a deterministic traversal of the attribution graph. The process includes loading the relevant parent relationships, computing contribution weights, routing each share to the appropriate vault, aggregating small transfers when possible, and executing settlement in a single atomic transaction.

This workflow resolves complex royalty structures without manual intervention and prevents intermediate inconsistencies by ensuring that each distribution either completes entirely or is reverted in full.

5.3 Vault Architecture

Each IPNFT is associated with a Royalty Vault that serves as the canonical collection and payout point for earnings. Vaults enforce predefined percentage splits, support minimum withdrawal thresholds to reduce the cost of small transfers, and provide transparent accounting through publicly visible balances. Beneficiaries withdraw earnings through a deterministic claim function, callable at any time.

Vaults separate royalty settlement logic from IPNFTs, creating auditable revenue distribution infrastructure. Each Vault issues 1 billion fungible tokens to ensure sufficient granularity for fractional royalty ownership. Token holders call the vault to claim their proportional share of accumulated revenue, enabling liquid markets for IP royalty streams and programmable revenue splits among collaborators.

5.4 Optimization Strategies

Origin incorporates several strategies to achieve scalable royalty processing. Batch execution reduces overhead by combining multiple royalty events into a single settlement. Lazy evaluation postpones certain computations until a beneficiary initiates a withdrawal. State compression stores only essential lineage data on chain. Graph caching accelerates repeated traversals for frequently accessed attribution paths.

Together, these optimizations significantly reduce gas consumption compared to naive implementations and make royalty distribution viable for high-volume derivative ecosystems.

6 Agent Integration

6.1 Agents as IP Entities

Origin treats AI agents as full participants in the IP economy. An agent can register its outputs as IPNFTs, obtain licenses for datasets or models used in training or inference, negotiate terms through predefined strategies, and distribute royalties to upstream contributors. This design allows agents to operate as autonomous economic actors while preserving transparent ownership, lineage, and compensation flows across the system.

6.2 Training Data Attribution

When an agent trains or fine-tunes a model using licensed datasets or other IPNFT-backed assets, Origin records the relationship between the training materials and the resulting output. The attribution process includes license validation for the input assets, calculation of relative contribution weights, creation of parent links in the new IPNFT, configuration of royalty routing rules, and documentation of provenance in both the metadata and the attribution graph.

These steps ensure that upstream data providers receive compensation when derivative models produce economic value. They also allow downstream users and auditors to verify how a model was trained and which assets contributed to its final form.

6.3 Autonomous Licensing

Agents can execute licensing decisions programmatically through smart contract interfaces. Decision engines incorporate market pricing, demand signals, usage context, and exclusivity requirements. Agents may adjust their licensing strategies based on historical behavior of counterparties, projected future value, or the competitive landscape within a particular application domain.

By using programmable licensing policies, agents can reliably participate in the IP economy without continuous human supervision. This supports a scalable autonomous ecosystem in which models, datasets, and software components engage in economic transactions with other agents while adhering to Origin's attribution and royalty rules.

6.4 IP Agentic Payments and x402

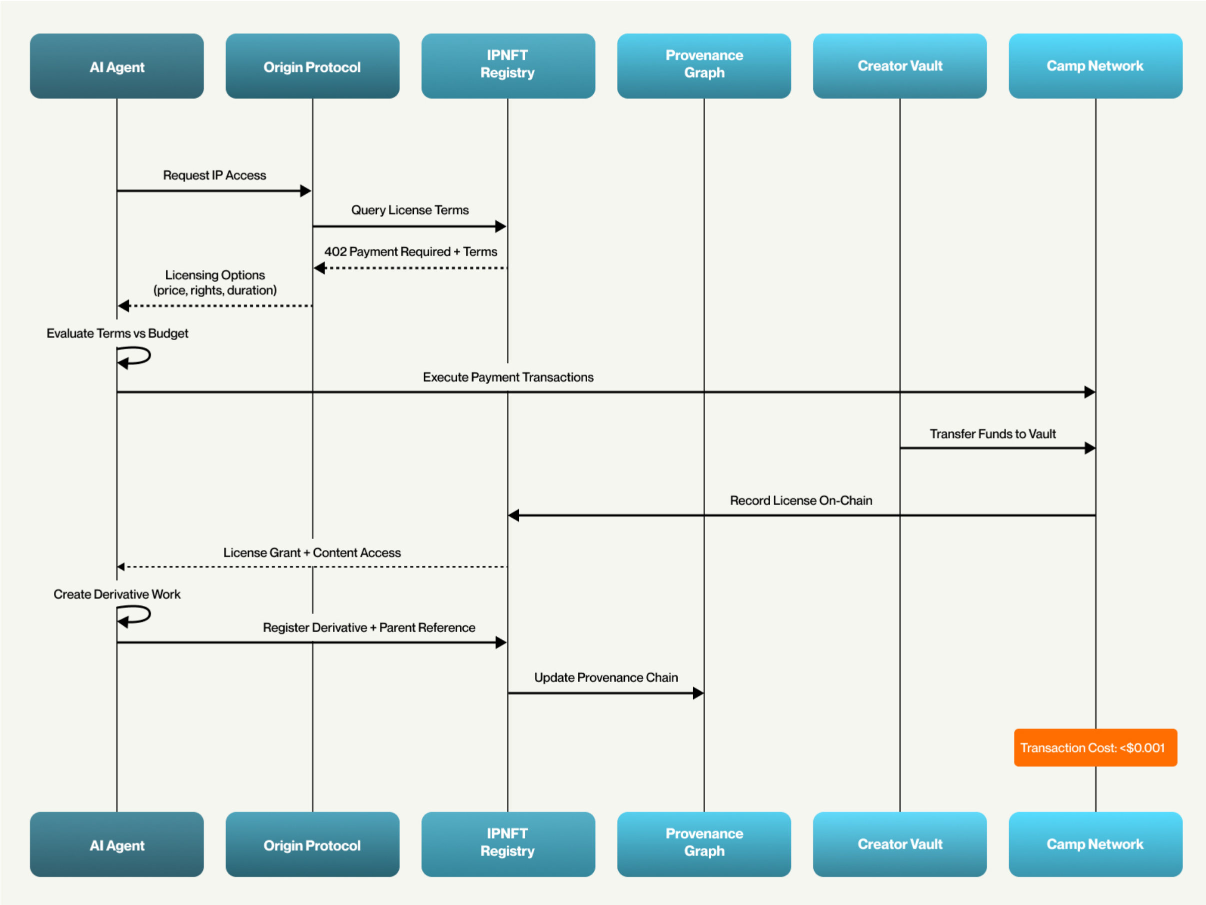

Figure 5: Agentic Payment Workflow

Origin enables AI agents to autonomously discover, license, and access intellectual property through x402 payment protocols. When an AI agent encounters protected content, the system responds with an HTTP 402 Payment Required status that includes machine-readable licensing terms: price, usage rights, attribution requirements, and the IPNFT identifier. This creates a permissionless marketplace for IP where agents operate as independent economic actors. An image generation model needing reference material queries the provenance graph for relevant IPNFTs, receives x402 responses with licensing options (commercial use, non-commercial, time-limited), selects appropriate licenses, and completes payment in milliseconds. The transaction is recorded onchain, establishing both legal rights and cryptographic proof of licensing. The agent then accesses the content and registers its derivative work in the provenance graph, automatically propagating royalties back to source creators.

7 Gateway Services

7.1 Content Delivery Gateway

Origin's gateway layer enables authenticated access to licensed content while maintaining decentralized permission verification. The gateway performs three core operations: (1) validates license ownership through cryptographic proofs against the onchain registry, (2) confirms the license grants appropriate rights, and (3) facilitates secure content delivery via encrypted streaming or controlled download.

The gateway also maintains usage accounting to support analytics, royalty calculations, and compliance auditing. By centralizing content delivery while decentralizing permission verification, the gateway abstracts storage complexity from applications and ensures all access decisions remain anchored to immutable onchain license state.

7.2 Legal Document Bridge

Commercial and exclusive licenses require legally enforceable contracts, recognized by traditional legal systems. Origin's legal bridge automates generation of jurisdiction-compliant agreements that reference onchain state as their authoritative source.

The bridge generates jurisdiction-specific contract templates, embeds license terms directly from blockchain records (license ID, payment schedules, territorial restrictions, duration), facilitates digital signature collection with optional notarization, and stores cryptographic hashes of executed agreements onchain with complete documents on IPFS.

Each contract explicitly references its corresponding onchain license record, creating bidirectional verification: the legal document cites the blockchain transaction as the definitive statement of terms, while the onchain record includes the IPFS hash of the executed agreement. This ensures licenses retain legal enforceability in traditional courts while leveraging blockchain transparency for evidence and dispute resolution.

7.3 Dispute Resolution Framework

Origin implements an on-chain dispute resolution system that combines economic incentives with community governance to address IP infringement and fraudulent claims.

Dispute Initiation: Any community member can raise a dispute against a potentially infringing or fraudulent IP asset by calling the `raiseDispute` function, submitting evidence alongside a fixed dispute bond in $wCAMP. This economic stake ensures good-faith participation and deters frivolous claims.

Assertion Period: Once a dispute is raised, the IP owner has a fixed dispute cooldown period to assert their position by calling the `disputeAssertion` function and providing counterevidence. During this cooldown period, the dispute initiator retains the right to cancel their claim by calling the `cancelDispute` function, provided the dispute has not yet been asserted. If the IP owner fails to respond within the cooldown window, the contract owner marks the dispute as `Valid`, automatically resolving it in favour of the initiator.

Community Adjudication: When an IP owner successfully asserts a dispute, the protocol enters the `disputeJudgementPeriod`. During this window, active $CAMP stakers exercise their governance rights by voting on the resolution. Once the judgment period concludes, the dispute is resolved based on the voting outcome, ensuring decisions reflect community consensus and the economic alignment of token holders.

Cascading Disputes: To maintain integrity across the IP graph, any community member may call the `tagChildIp` function to mark child IPs as disputed when a parent asset has been declared disputed. This mechanism ensures that derivative works are appropriately flagged during ongoing disputes, preventing the propagation of potentially infringing content through the attribution graph.

8 Economic Model

8.1 Fee Structure

Origin applies a fee model designed to sustain long-term protocol operation without imposing excessive costs on participants. Registration fees cover the creation and validation of new IPN-FTs. License fees apply to transactions involving access rights or commercial use. Royalty fees represent a small share of distributed royalties that support protocol infrastructure. Gateway fees use a usage-based model for secure content delivery and bandwidth consumption.

These fees are calibrated to balance economic sustainability with competitive pricing relative to traditional IP management systems.

8.2 Incentive Alignment

The economic design aligns incentives across creators, users, validators, and developers. Creators benefit from automated licensing flows, transparent attribution, and reliable royalty settlement. Licensees gain instant access to clear pricing, predictable terms, and machine-verifiable rights. Validators receive a portion of protocol fees for maintaining security and system integrity. Developers are supported through grants and incentives that encourage the creation of new applications and integrations.

This alignment encourages active participation and supports steady ecosystem growth.

9 Security Analysis

9.1 Threat Model

Origin evaluates security across economic, technical, and social dimensions. Economic risks include Sybil-based registrations, attempts to manipulate royalty flows, and extraction of undue fees. Technical risks involve smart contract vulnerabilities, unintended state transitions, and denial of service scenarios. Social risks include identity misuse, intentional plagiarism, and incorrect attribution claims.

The threat model defines these vectors so that each can be addressed through targeted protocol mechanisms.

9.2 Mitigation Strategies

Origin employs a defense-in-depth approach that combines economic, technical, and social safeguards. Economic protections include registration bonds, slashing conditions for malicious behavior, and rate limits to prevent abuse. Technical protections include formal verification of critical components, structured upgrade paths, and circuit breakers that pause system operations under abnormal conditions. Social protections include optional identity verification, reputation mechanisms for creators and agents, and a multi-tier dispute resolution process.

These measures collectively reduce attack surfaces and reinforce the integrity of the IP ecosystem.

9.3 Audit and Verification

The protocol undergoes continuous security assessment through both internal and external processes. Independent firms audit smart contracts before deployment and after major revisions. Critical modules are subject to formal verification to prove correctness under defined assumptions.

This comprehensive validation framework supports the long-term reliability and safety of the Origin protocol.

10 Implementation Roadmap

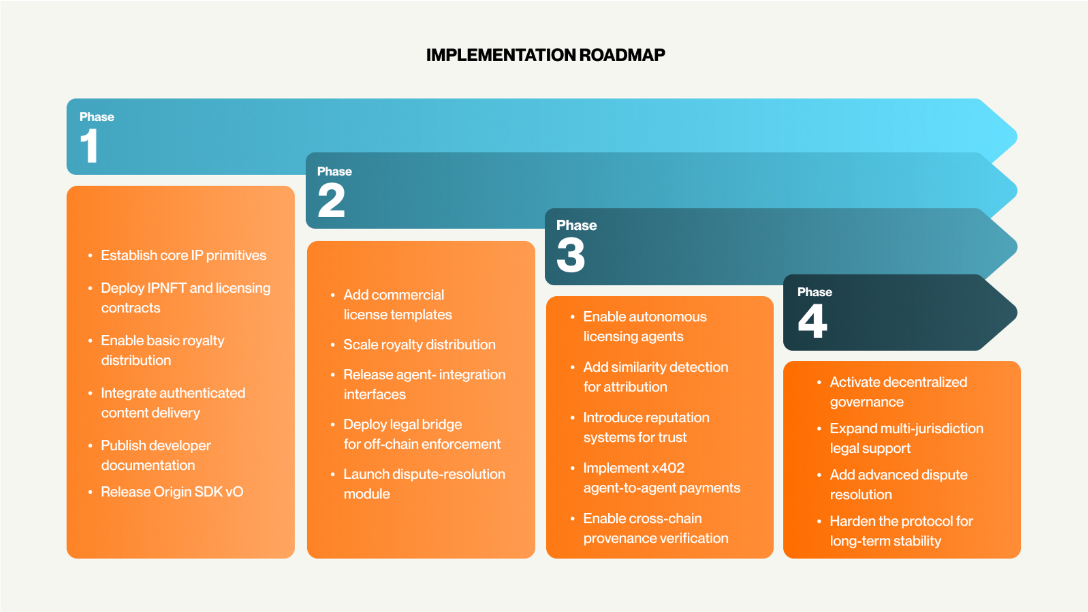

Figure 6: Origin Roadmap

Phase 1: Foundation (Q1 2026)

The initial phase focuses on establishing the core technical primitives required for IP registration and licensing. This includes deployment of the IPNFT and licensing contracts, implementation of basic royalty distribution, integration of authenticated content delivery through the gateway layer, and publication of developer documentation to support early builders.

Phase 2: Expansion (Q2 2026)

The second phase extends the protocol's economic and legal capabilities. Commercial license templates are introduced, royalty distribution is optimized for large-scale use, and standardized interfaces for agent integration are released. The legal bridge is deployed to support enforceable off-chain agreements.

Phase 3: Intelligence (Q3 2026)

The third phase introduces intelligence-driven features. Autonomous licensing agents manage pricing and negotiation, similarity detection algorithms support attribution and originality checks, reputation systems enhance trust among creators and autonomous participants, and cross-chain verification enables interoperable provenance.

Phase 4: Maturation (Q4 2026)

The final phase focuses on hardening and governance. Decentralized governance mechanisms are activated, multi-jurisdiction legal frameworks are expanded, advanced dispute resolution processes are introduced, and the protocol begins to stabilize into a mature and stable state suitable for long-term reliability.

This phased rollout ensures the protocol evolves predictably while maintaining compatibility across all deployed components.

11 Conclusion

Intellectual property management has historically relied on static contracts, manual verification, and enforcement through litigation structures built for a slower, human-mediated world. These systems cannot support digital environments where content circulates at algorithmic speed and AI agents operate continuously as creators, licensors, and consumers of IP. Origin solves this gap by rendering rights machine-readable and cryptographically enforceable, turning licensing into a programmable onchain primitive executed with deterministic guarantees.

The protocol delivers capabilities traditional IP frameworks cannot: composable licensing logic, hybrid onchain and legal enforcement, agent-native execution paths, transparent royalty flows, and cryptographic provenance. Together, these enable new forms of creative collaboration and micro-licensing economies that function at the speed and scale of modern AI systems.

As autonomous agents take on a greater share of content generation and distribution, they will query registries, acquire rights, create derivatives, and resolve disputes programmatically without human mediation. This infrastructure becomes essential as creative markets accelerate toward billions of daily micro-transactions. Challenges remain: cross-jurisdictional enforcement, registry scalability, governance, and evolving AI policies. However, Origin's modular architecture is designed to incorporate these constraints iteratively rather than attempt a monolithic solution.

Digital intellectual property requires digital infrastructure. Origin demonstrates that programmable, autonomous rights management is not theoretical but operational today, with future progress driven by ecosystem adoption, regulatory alignment, and continuous refinement.

12 References

1. Hadfield, G. K., & Koh, A. (2025). *An Economy of AI Agents*. NBER Handbook on the Economics of Transformative AI. arXiv:2509.01063. Retrieved from https://arxiv.org/abs/2509.01063

2. Makridis, C., Bates, T., Barnes, N., & Glasgow George, J. (2024). *The Future of Intellectual Property: The Complementary Role of Artificial Intelligence and Blockchain*. Journal of

the Patent and Trademark Office Society, 104:293.

3. Camp Network. (2025). *Camp: The Sovereign Layer 1 for Creative Intelligence*. Technical Whitepaper.

4. Ethereum Foundation. (2018). *ERC-721: Non-Fungible Token Standard*. Ethereum Improvement Proposals. Retrieved from https://eips.ethereum.org/EIPS/eip-721

5. Ethereum Foundation. (2023). *ERC-6551: Non-fungible Token Bound Accounts*. Ethereum Improvement Proposals. Retrieved from https://eips.ethereum.org/EIPS/eip-6551

6. Creative Commons. (2024). *Creative Commons License Framework*. Legal Documentation. Retrieved from https://creativecommons.org/

7. World Intellectual Property Organization. (1996). *WIPO Copyright Treaty*. International Treaty. Retrieved from https://www.wipo.int/treaties/en/ip/wct/

8. x402 Foundation. (2025). *x402: A Payment Protocol for the Internet*. Technical Whitepaper. Retrieved from https://x402.org/