



TRUST: Intuition's ERC-20 Token for Decentralized Data Creation, Curation, Ownership, and Discovery

Intuition Systems, Inc.

Abstract

The accelerating proliferation of machine- and human-generated content has outpaced society's ability to evaluate authenticity and relevance, threatening an epistemic commons dominated by low-cost duplication and strategically amplified falsehoods, and the value captured from this data has centralized into the hands of the few. We therefore introduce TRUST-Intuition's native ERC-20 governance and utility token-a cryptoeconomic primitive that democratizes access to the monetization of the world's data, puts ownership back into the hands of the people, and assigns explicit financial weight to every claim recorded in the network's token-curated knowledge graph. All write and read operations are metered in TRUST: authors place stake when publishing, and consumers pay 'gas fees' when querying. These requirements deter sybil proliferation and denial-of-service spam while ensuring that every interaction carries auditable economic gravity. Backed by TRUST, data becomes a tradeable on-chain asset that can be permissionlessly owned by anyone, whose market value can appreciate or decay as evidence, sentiment, and usage evolves. The endogenous price signal continuously ranks information, rewarding agents who underwrite verifiable data and imposing measurable costs on those who subsidize error and irrelevancy. Beyond curation, TRUST is also bonded by independent node operators who index and serve the graph, earning rewards and facing slashing for downtime or tampering. Finally, the same tokens confer governance rights: holders propose and vote on protocol upgrades, fee parameters, bonding-curve templates, and treasury allocations, ensuring that the rules of the network evolve under the control of the very community that secures and curates its knowledge.

The mechanism integrates (i) programmable bonding curves that seed liquidity around nascent claims, (ii) time-weighted bonding that amplifies long-horizon conviction, (iii) fee-redistribution pipelines that recycle revenues to creators and curators, and (iv) network

incentives that stream emissions to operators who maintain the network and faithfully serve graph data. An on-chain governance layer lets token holders tune decay constants, bond sizes, slashing curves, reward weights, and emission schedules, keeping all incentive gradients adaptive to network dynamics.

We thereby furnish a scalable, permissionless blueprint for a token-weighted Semantic Web of Trust in which credibility is continuously collateralized rather than retrospectively adjudicated. Empirically, this architecture channels capital toward durable signal, accelerates the liquidation of misinformation, and internalizes the social cost of epistemic failure—transforming passive content storage into an economically self-refining substrate for collective intelligence.

Introduction

Digital content now grows hyper-exponentially, driven by frictionless creation tools, large-language-model paraphrasing, and algorithmic amplification. The traditional mechanisms that once filtered signal from noise—editorial review, domain expertise, and slow reputational build-up—cannot keep pace with this production frontier. Instead, contemporary discovery has defaulted to cost-free social gestures: likes, shares, reposts, and comments. Although these gestures constitute a vast, continuous act of curation, their economic yield accrues almost entirely to the Web 2 intermediaries that harvest the resulting engagement data for advertising and attention arbitrage. End-users, who supply the curation labour, neither own the informational assets they elevate nor participate in the revenue they generate.

A parallel domain—financial markets—demonstrates a more robust discovery paradigm. A security's price and liquidity simultaneously broadcast relevance and allocate capital: liquid, high-valuation equities automatically appear in indices, analyst coverage, and institutional portfolios. Visibility and value reinforce each other in a reflexive loop that is transparent, auditable, and economically settled. Crucially, participants who supply price-discovery services are compensated through spread capture, alpha, or dividend flow.

Intuition generalises this market-based discovery loop to all data. At its core lies an open, token-curated knowledge graph in which every claim, identity, or semantic relationship is associated with a financial primitive denominated in **TRUST**. A positive stake functions as a priced endorsement; a counter-stake operates as a short position against accuracy or

relevance. The aggregate depth of these positions defines a continuously updated "price curve" for epistemic weight, recasting the knowledge graph as a live order book for truthfulness and utility.

Formally, let C_i denote a claim, with s_i^+ and s_i^- representing the outstanding long and short **TRUST** positions, respectively. The net relevance weight

$$w_i = s_i^+ - \kappa s_i^-, \kappa \in (0, 1]$$

serves both as a ranking signal and as a fee-distribution coefficient: query revenues and bonding-curve spreads flow to addresses in proportion to their stake-adjusted contribution to w_i . The weight of dormant positions decays via a governed half-life parameter τ , ensuring that relevance is continuously re-priced as new information and adversarial challenges arrive.

By embedding this Token-Weighted Web of Trust, Intuition achieves four objectives:

1. **Economic Skin-in-the-Game:** Endorsements and refutations carry explicit opportunity cost, discouraging sybil amplification and cheap manipulation.
2. **Programmable Revenue Sharing:** Curators and validators receive pro-rata fee flows, rectifying the value-extraction asymmetry characteristic of Web 2 platforms.
3. **Self-Cleaning Knowledge Base:** Automated weight decay and counter-position profitability impose a persistent cost on misinformation, driving spurious data toward economic liquidation.
4. **Community-Governed Parameters:** Holders of **TRUST** vote on decay constants, bonding-curve shapes, and fee splits, fine-tuning incentives as the network evolves.

The remainder of this whitepaper elaborates on the cryptoeconomic architecture that supports these properties: supply schedules, staking and slashing logic, governance mechanics, and network operation. Taken together, they constitute a scalable blueprint for coupling market dynamics with knowledge curation—rewarding the act of producing and validating useful information while imposing measurable costs on irrelevance and fraud.

Token Design

TRUST is an ERC-20 compliant token that serves as the native asset of the Intuition network. Its design balances **utility** and **governance** functions to drive network effects in the

knowledge graph while ensuring the community can steer the platform's evolution. In total, the token supply is planned with a combination of initial allocation and ongoing emissions, incentivizing early adopters and ensuring sufficient supply for future contributors.

Utility and Value: The **TRUST** token underpins economic transactions and incentives within Intuition. Core creating, reading, and staking actions on data items use the Intuition's native token, TRUST. In this manner, **TRUST** acts as the unit of account for rewards and a measure of one's contribution to the system. Users earn **TRUST** through productive actions (adding high-quality data, staking early on accurate claims, contributing to governance, etc.), effectively gaining **equity in the data** they help curate. This tokenized "data equity" entitles contributors to ongoing value as the network grows – for example, when a piece of information becomes widely referenced, those who curated it receive a share of fees or token rewards proportional to their stake and timing of contribution. Conversely, bonding **TRUST** also grants access to platform privileges, such as participating in governance votes and potentially receiving a portion of global platform yield, as described in later sections.

In addition to incentivizing data curators, the Intuition token model explicitly incorporates network **node operators** as a core participant in the ecosystem's cryptoeconomic design. Node operators **must bond a certain amount** of the TRUST token as collateral to run an Intuition indexing node which maintains and serves the Intuition knowledge graph. At network launch this bonding requirement is planned to be a flat (fixed) amount of TRUST for every operator, ensuring a uniform stake and lowering complexity for new participants. This bonded stake gives each operator significant "skin in the game," aligning their interests with the network's health and security. While initially flat, the bonding model is designed to be adaptive via governance – the community can later vote to adjust it into a dynamic requirement based on operator activity or network needs. For example, in the future the bond might scale with the volume of indexed claims or queries an operator serves, ensuring that those handling more data (and earning more fees) put up commensurate collateral. This flexibility in the token design allows Intuition to calibrate security as the network grows, all while keeping token holder governance at the center of these economic adjustments.

Crucially, the bonded stake held by node operators is subject to slashing penalties for misconduct or underperformance. Slashing is tiered by severity to deter bad behavior at multiple levels. Minor lapses – such as failing to meet the minimum uptime threshold (e.g. not maintaining $\geq 95\%$ availability over a given epoch) – trigger a partial slashing of the operator's staked TRUST. In this case, a percentage of the bond is forfeited as a penalty for downtime,

providing a strong incentive to maintain robust and reliable infrastructure. By contrast, any egregious violations of protocol honesty result in full slashing of the bonded stake. If a node operator is proven to have tampered with data or served falsified (incorrect) knowledge graph responses, the entirety of their bonded TRUST can be confiscated, effectively ejecting them from the network. These slashing mechanisms make it economically irrational to behave maliciously: an operator stands to lose substantial value for attempting to deceive or for neglecting their duties. By requiring bonded tokens and enforcing slashing, the Token Design ensures that node operators are trust-aligned stakeholders who only profit by keeping the network truthful and highly available. This extension of the TRUST token's utility – from curating information to also securing its availability – strengthens the overall cryptoeconomic model of Intuition.

Token Supply and Emission: The token's supply model is designed to **reward early and long-term participants** while maintaining a sustainable inflation rate. A portion of **TRUST**'s supply is set aside for initial distribution, ensuring the community of early Intuition adopters is recognized. The remainder of tokens are allocated to a **minting schedule** over the coming years as "emissions" to be earned by the community. The emission schedule follows a decaying curve to favor early network growth yet still provide incentives in the long run. For example, if T_0 is the initial emission rate, the rate at year t could follow a geometric decay:

$$E(t) = E(0) \times (1 - \rho)^t$$

where ρ is a yearly decay factor (e.g. $\rho = 0.20$ for a 20% annual reduction). This ensures that the largest token rewards are available during the network's bootstrapping phase, while a long "tail emission" provides continuous incentives for ongoing contributions and securing of the knowledge graph. The **total supply** of **TRUST** may be capped with a hard cap reached asymptotically under the emission schedule, or governed by token holders who could adjust ρ or halt emissions via governance if needed.

This allocation ensures that a sufficient supply of tokens are earned by the community over time rather than given upfront, aligning with **TRUST**'s role as a participation-driven token. Team and investor allocations are intentionally moderate and vesting-protected to prevent excessive centralized control, and a meaningful portion is dedicated to the community from the start.

In summary, **TRUST**'s design treats *knowledge as an asset*: those who contribute to Intuition's knowledge graph earn tokens which entitle them to shared **ownership of that**

knowledge's future value. As more users and applications leverage Intuition's data (for example, an AI agent querying the graph or a dApp integrating Intuition for identity verification), the demand for accurate information – and thus the demand for signaling that accuracy via staking – increases. This creates a virtuous cycle where **TRUST** captures a portion of the value generated by the network effects of reliable data. The token's governance properties (discussed later) further ensure that those who hold **TRUST** have a voice in refining the rules of the system, completing the feedback loop between **token value** and **network health**.

Utility Mechanisms: Data Creation, Staking, Signaling, & Node Operation

One of the primary innovations of Intuition is how it uses tokens and staking to curate a high-quality, **trustworthy knowledge graph**. **TRUST** plays a pivotal role in four interrelated utility mechanisms: **data creation**, **staking/signaling**, **querying**, and **network maintenance**. Together, these mechanisms incentivize users to add and serve truthful information and align on common standards, while penalizing misinformation or low-quality contributions through opportunity cost.

Data Creation Incentives

In Intuition, anyone can create a new **Identity** (called an *Atom* at the protocol level) or a new **Claim** (a semantic linkage of atoms, also called a *Triple*). For example, a user might create an identity for a new concept or entity, or assert a claim like "[Restaurant X] [isLocatedIn] [City Y]". Data creation is permissionless, but not free of cost – this is by design. Whenever a user introduces a new identity or claim, a small **creation fee** (denominated in Intuition's native **TRUST** token) is required. This fee may be likened to a 'gas cost' in traditional blockchain networks, and serves two purposes: it prevents spam by making large-scale fake data submissions economically costly, and it seeds the economic value of the new data item. A portion of the creation fee is converted into an initial "stake" on that data, effectively giving the creator a starting **ownership stake** in the claim or identity they just added. In other words, by paying the fee the creator is also **staking** on the truth or relevance of that data point. They signal "I vouch for this information" with economic weight. Immediately, this stake is

represented in the data's **vault** (each Atom or Triple has its own escrowed stake pool, as described below) and grants the creator some **stakeholder shares** of that item. The initial creator may stake any amount of TRUST above and beyond the initial creation cost to increase their exposure to the claim without risk of fruntrunning.

In return for creating valuable data early, the user is positioned to earn future rewards: if the identity or claim they added becomes widely used and trusted, the creator's stake shares grants them a portion of fees and token emissions related to that item. This implements the notion that **early contributors earn equity in the data** they help introduce. For example, adding a widely-used predicate like `[hasReview]` early on could yield significant **TRUST** rewards over time as others use that predicate in numerous claims. This mechanism incentivizes users to add information they genuinely believe will be useful and true (since they invest a fee and want returns), aligning individual profit motive with the accuracy and utility of the shared knowledge graph.

Staking Mechanism on Identities (Atoms) and Claims (Triples)

Staking is the process by which any user can deposit value to **signal support or opposition** for a given identity or claim. In Intuition's architecture, each Identity (Atom) and Claim (Triple) is associated with one or more **vaults** – typically a single vault for an Identity, and two vaults for a Claim: one "True" vault for support of a Claim, and one "False" vault for the rejection of a Claim. Users stake by depositing **TRUST** into these vaults to express their conviction about the data:

- Staking on an Identity implies you consider that identity (the entity or concept) important or valid in the context of the knowledge graph. High stake on an identity signals to the network that this entity is worth attention and likely to be referenced elsewhere.
- Staking on a Claim in the positive vault means you endorse the claim as **true** or useful; staking in the negative vault means you dispute its validity. This many-to-one staking model allows multiple independent users to weigh in on the truth of a single claim.

Staking confers **multiple benefits**:

1. **Signal of Trust:** The amount staked serves as a quantitative signal to others about the credibility of the information. Much like upvotes or ratings in Web2 systems, stake

represents belief – but backed by real value, which makes signals costly to fake.

Information with more stake behind it will be ranked higher or considered more canonical in the knowledge graph, helping consumers find reliable data quickly.

2. **Fee and Reward Sharing:** Stakers effectively become shareholders of that data item's "micro economy." They are entitled to a share of any fees that future interactions with that data generate (described under Platform Revenue). For instance, if a claim accumulates significant stake and usage, new contributors who stake later or reference that claim will pay fees that get distributed to existing stakers as rewards. Early and sizable stakers thus earn a return on their commitment, proportional to their share of the total stake.
3. **Token Emissions:** A portion of **TRUST** emissions may be allocated to active stakes on identities and claims. The protocol can periodically reward those who have staked on data with freshly minted **TRUST**, as an incentive for curating the graph. For example, if 1000 **TRUST** are emitted per day to the staking program, they might be divided among all vaults according to some formula (e.g., proportional to the square root of total stake, to favor breadth) and then further split among stakers in each vault by their fraction of that vault. This encourages users to stake on a wide range of data that they believe will become valuable, not just chase one popular item.
4. **Governance Participation:** In some designs, staking on content could double as a lightweight governance signal for that content's metadata or status. However, in Intuition, content staking is separate from general protocol governance. Still, staking on data indirectly influences governance by setting de facto standards (explained next) – the community might formalize heavily staked standards into official defaults.

When a user stakes amount $S_{i,j}$ on a claim j , they receive **shares** $sh_{i,j}$ representing their proportion of ownership in that claim's vault. If the total stake in claim j 's True vault is $S_{\text{total},j}$ after their deposit, then the fraction of stake they own is:

$$\text{StakeFraction}_{i,j} = \frac{S_{i,j}}{S_{\text{total},j}}.$$

These shares entitle the user to the same fraction of any rewards distributed to that vault, whether those are fee revenues or token emissions. Unstaking (withdrawing stake) is allowed, but are subject to an exit fee to prevent abuse like rapid stake flipping. If a user exits, they give up future rewards from that item, which will then be enjoyed by the remaining stakers.

It's important to note that **staking is dynamic**: as new information emerges or opinions change, users can rebalance their staked positions. Intuition's design allows a many-to-one

relationship between attestations and claims – meaning many people can stake on (attest to) the same claim, and a single user can stake on multiple claims. This flexibility supports a **crowdsourced consensus** model: truth in Intuition is determined not by any central authority, but by where users collectively allocate their economic weight.

Bonding Curves in Vaults and Governance

Bonding curves play a central role in Intuition’s Vault mechanism for Atoms and Triples, defining how staking is priced and rewarded. Each Vault is essentially an automated market maker specific to an Atom or Triple, where the price of “shares” (staking positions) is determined by a curve rather than an order book. This means as more users stake on a given piece of data, the cost to acquire additional stake adjusts algorithmically. By utilizing bonding curves, early supporters of data that prove to be *popular and true* are economically rewarded (they obtain stake shares cheaply and can later redeem at higher value). Conversely, latecomers pay higher prices for the same stake, reflecting the increased confidence or consensus in that Atom or Triple. In this way, bonding curves turn each Vault into a **liquidity-free market** for truth signals – anyone can stake or withdraw at any time without needing a buyer or seller on the other side, since the Vault’s smart contract continuously prices the shares based on the curve. This approach is crucial for scaling to potentially *billions of claims*, as it eliminates the need for manual market-making on each claim and ensures continuous liquidity and price discovery within each Vault.

Bonding Curve Registry: Intuition allows Vaults to be associated with any number of bonding curve formulas via a *Bonding Curve Registry*. All Atoms, Triples, and Counter-Triples start with a default **Pro Rata (Linear) curve**, which maintains a constant price-per-share ratio. This linear curve is the safest option – it keeps staking and redeeming risk low and predictable (you generally get out what you put in, aside from minor fees). However, users can opt into more dynamic bonding curves when depositing or withdrawing from a Vault. The registry provides a library of curve shapes that any Vault can use, each implemented through a common interface for converting between deposited assets and share price. The Intuition DAO governs this registry, curating which curve types are available and configuring their parameters as needed. By voting to add new curve formulas or adjust existing ones, the DAO can introduce new incentive structures for staking behavior in a controlled, transparent way. This governance ensures that as the network grows and new data types emerge, the available bonding curves can evolve accordingly (always via community consent).

Bonding Curve Shapes and Staking Incentives

Different bonding curve shapes instantiate different economic “games” for how users stake on Vaults. The shape of the curve dictates how strongly early staking is rewarded versus how expensive late staking becomes. Below we describe several key bonding curve types in the registry and give examples of what kinds of data or contexts they are best suited for:

Linear (Pro Rata) Curve: A linear bonding curve increases the token price at a constant rate as supply (staked volume) grows. In formula terms, it can be expressed as $P(\text{supply}) = m \cdot \text{supply} + b$ (a straight line). This is the default **Pro Rata** curve in Intuition’s Vaults. It neither heavily favors early nor late participants – every additional stake increases the share price by the same fixed amount. As a result, linear curves are **low-risk and predictable**. They are ideal for data where a steady signal is desired without speculative volatility. For example, an Atom representing a stable concept or identity might use a linear curve so that stakeholders can enter and exit without dramatic price swings. Linear Vaults primarily serve to aggregate confidence on an identity or claim while minimizing gamification; early stakers don’t gain an outsized advantage, but they also aren’t exposed to large downside risk. This fairness makes linear curves a good fit for foundational data that benefits from broad, ongoing participation.

Exponential (Power Law) Curve: An exponential bonding curve makes the price increase *accelerate* as more tokens are minted. In other words, if the staked supply doubles, the price **more than doubles**. A simple example is a quadratic curve $P(\text{supply}) = a \cdot \text{supply}^2$, which is a special case of exponential growth. Exponential curves strongly **reward early stakers**: when you stake on a Vault early (when supply is low), you pay a very low price per share. As the Vault’s supply grows with later users, the price per share skyrockets, meaning your early shares can now be sold back at a much higher price. This mechanism is well-suited for **rapidly emerging or trending data** – for instance, a newly created Triple (claim) about a breaking news topic. Early believers in the claim take on risk but could profit greatly if many others subsequently find the claim important (driving the stake price up sharply). Projects aiming to encourage **bootstrapping of truth discovery** often consider power-law curves, since they create strong incentives to be the first to stake on information that later turns out to be valuable. However, exponential curves can also lead to volatility; if a trend reverses and people withdraw, the price falls quickly on the way down as well (the mirror of the rapid rise). The DAO will typically vet exponential curve parameters (the base and exponent) to ensure they’re not too extreme for the intended context, keeping the game dynamic but not outright chaotic.

Sigmoid (Logistic) Curve: The sigmoid curve (S-shaped logistic function) offers a **balanced, multi-phase incentive** structure. At first, when supply is low, the curve grows slowly (nearly flat pricing) to encourage participation – similar to linear or even a constant price for initial tokens. As confidence builds and the stake supply enters a growth phase, the price ramps up **steeply** in the mid-section of the curve (the inflection point). This creates a window where early adopters who got in cheap see significant gains as the token price surges. Eventually, the curve **flattens out** again at high supply, meaning once a claim or Atom has a very large backing (widely accepted as “truth”), additional staking doesn’t increase the price as drastically. Latecomers are “heavily charged” in the sense that by the time the curve has flattened at the top, the price per share is quite high – but importantly, it’s capped in growth rate, avoiding infinite runaway prices. Logistic bonding curves are well-suited for data that you expect to **go viral or reach widespread adoption**, following a typical adoption lifecycle. For example, consider a Vault for a claim that a certain technology will be adopted globally: initially, only a few visionaries stake (price low), then a critical mass arrives and price shoots up, finally almost everyone agrees on it and the price plateaus. From a governance perspective, sigmoid curves are useful when the DAO wants to encourage early staking but also wants a **natural leveling-off** as consensus is reached, to avoid late-stage speculation. The curve can be defined by a formula such as $P(s) = \frac{L}{1 + e^{-k(s-s_0)}}$, where L is the upper price limit (plateau), and s_0 is the supply at the inflection point. By adjusting these parameters in the registry, the community can fine-tune how quickly the “middle phase” of rapid price increase kicks in and what the maximum price asymptote should be.

DAO Governance and Curve Evolution

All bonding curves in the Intuition system are **DAO-curved**. Through the Bonding Curve Registry, the Intuition DAO can vote to add new curve types or adjust their parameters as new needs arise. Different data types benefit from different curve shapes, so this flexibility is critical. For example, factual data that becomes universally agreed-on might migrate toward a flatter or sigmoid curve as it matures, whereas fast-moving cultural trends might warrant sharper exponential curves to maximize early discovery. The governance process ensures that any addition of a curve (say, a new quadratic variant or a piecewise custom curve) goes through review – aligning the curve’s game mechanics with the **integrity and incentives of the platform**. Intuition’s stakeholders can propose curves optimized for particular domains (science facts, social claims, identity verification, etc.), and once approved, those curves become available for any Vault to use.

By empowering the community to configure bonding curves, Intuition creates a living economic system for truth-finding. **Curve shapes = incentive landscapes:** the DAO effectively decides what kinds of staking behaviors to encourage. Because every curve follows a standard interface in the contracts, adding a new one is seamless – the Vaults of any Atom or Triple can immediately start using the new curve to guide staking. This modular design means the platform can support a vast diversity of economic scenarios, all under one roof. Crucially, regardless of shape, all bonding curves maintain the property of **continuous liquidity** and algorithmic price discovery. Stakeholders can always enter or exit a position on a claim, and the price they pay or receive reflects the consensus and demand at that moment, as encoded by the curve's formula. This liquidity-free market design is what allows Intuition to scale to **billions of claims** without relying on external buyers or sellers for each piece of data – the bonding curves themselves handle the market dynamics. In summary, bonding curves in Vault staking provide the economic backbone for Intuition's governance of truth, enabling fine-tuned incentives, DAO oversight of game mechanics, and a path to global scale in decentralized knowledge curation.

Signaling and Token-Curated Knowledge

Staking is not a passive investment; it is an active **signaling mechanism**. By choosing where to stake, users effectively **vote with their tokens** on what data is important and correct. This creates a **Token Curated Registry** (TCR) dynamic for Intuition's knowledge graph. In a classic TCR, token holders stake to curate a list of entries (often deciding if an entry should be included or not). In Intuition, the "lists" being curated are the set of identities and claims in the knowledge base, and their relative trust or importance ranking. Good data, backed by many tokens, rises to prominence; bad or unverified data, lacking economic support (or opposed via negative stakes), remains obscure or is flagged.

Signaling via stake has several profound effects on the Intuition ecosystem:

- **Consensus Identifiers for Everything:** In a permissionless system anyone can label the same concept in countless ways, replicating today's web-wide chaos of duplicate IDs. Intuition turns that problem into a coordination game: the more stake an identifier attracts, the more fees flow to everyone who uses it. Curators therefore gravitate toward the tag with the deepest economic gravity, and divergent labels quickly lose steam. When you publish data you may simply choose the identifier that already carries the highest

token-weighted signal—knowing it will interoperate with the rest of the graph and maximise your own rewards. Or, if you think you have a better identifier that can reach critical mass, you may opt to 'fork' and choose this identifier over the consensus identifier, hoping that it becomes the consensus identifier over time and thus rewarding you asymmetrically for having an early stake in a popular Atom/identifier. Over time this price-driven convergence yields consensus identifiers for every entity, eliminating costly reconciliation work and binding the knowledge graph into a coherent, structured whole.

- **Emergent Ontology Standards:** Over time, as users independently add data, certain ways of describing things will gain more traction. For example, imagine multiple predicates exist for the concept of a product review: `[hasReview]`, `[review]`, `[userOpinion]`, etc. If the community largely stakes behind `[hasReview]` when adding review data, this predicate accumulates the most support. The stake signals push it to become the *de facto* standard, while alternate, less-staked predicates fall out of use. Thus, Intuition achieves **cryptoeconomic consensus on data standards**: the network financially rewards alignment on common schemas, solving the fragmentation that plagued earlier attempts at a Semantic Web. Early adopters of what becomes the winning standard reap more rewards, which is a direct incentive to converge on the best way to represent data.
- **Data Quality and Truthfulness:** Stakes are signals that cost users real money (opportunity cost if nothing else), which discourages casual support of false or irrelevant claims. If a claim is contentious, ideally it will have stake on both the True and False side (Intuition supports **counter-staking** to signal disagreement). The magnitude of stake on each side signals the community's confidence in the claim's veracity. In extreme cases, governance mechanisms (or automated rules) might resolve contested claims by rejecting those with insufficient support or by slashing fraudulent stakes if an objective oracle is available. But in general, the presence of even modest opposing stake will warn users that a claim is disputed, thereby upholding **data transparency**. A rational actor will not stake heavily on a false claim because if the truth comes out (others stake on the opposite side or no one validates it further), their stake will not earn rewards, will result in reputational damage, and may even incur a loss (lack of future fees, or potential slash if implemented). The safest way to earn through staking is to find and back **truthful, useful information early**, which is exactly the behavior the platform seeks to encourage.
- **Identity/Claim Weight:** Staking also dictates how visible that identifier (or any individual claim) becomes inside the network. All **TRUST** staked for an identity or claim is a standing order to Intuition's network of nodes: "prioritise this datum." The larger and

longer-lived the deposit, the stronger the guarantee that the data will be fetched, cached, and served at low latency to downstream apps, agents, and analytics pipelines.

Conversely, claims with little or no stake sink toward archival storage and may be pruned from fast-access layers altogether. If you want your information to surface in wallets, explorers, and AI assistants, you (or other interested parties) must continuously underwrite it with TRUST. In effect, stake functions as bandwidth rent—paying for the network's attention and storage while signalling that the claim is valuable enough to deserve that real-time availability.

- **Reputation:** Curators who repeatedly back accurate, valuable claims become high-signal nodes in the network's Web-of-TRUST: other users (and autonomous agents) can subscribe to their staking feed, mirror their convictions, and route query fees through them. That following magnifies both influence and income: when your early bets inspire secondary stakes, you share in those downstream rewards or enjoy preferential fee splits as determined by governance. A solid reputation also grants soft power—your objections carry more weight in disputes, your proposed identifiers gain traction faster, and protocol upgrades may confer perks such as reduced bonding-curve friction or delegated voting rights. Lose credibility, however, and wallets, explorers, and AI assistants will down-rank or even ignore your assertions. In short, if you want the network to listen when you speak, cultivate a track record of surfacing useful truths—just like building trust and authority in the physical world.

In summary, **signaling** in Intuition transforms subjective assessments of truth and importance into quantifiable metrics on-chain. By staking tokens (economic weight), the community collectively curates the knowledge graph's content and standards. This innovative use of **TRUST** and staked value ensures that **the best information, not the loudest voices, rise to the top**, because backing a claim requires putting "skin in the game." The next sections will discuss how the **TRUST** token is further used to govern the broader parameters of this system and how its supply is managed to reinforce the incentives outlined here.

Emissions and Bonding

The **Emissions and Bonding** component of **TRUST**'s tokenomics deals with how new tokens enter the ecosystem and how participants can commit (or "bond") their tokens or other assets for enhanced rewards and features. A well-designed emission schedule encourages network

growth and fair distribution, while bonding mechanisms encourage long-term alignment by rewarding those who lock in their support for Intuition. This section describes the emission model for **TRUST** and introduces bonding mechanics, including formulas and tables to illustrate their effects.

Emission Model

As outlined in the Token Design section, **TRUST** token emissions serve to reward participants over time. The emission model can be characterized by a few key properties:

- **Decay Schedule:** Emissions start high and decrease over time (a high initial inflation that tapers). This is often implemented via a halving schedule or exponential decay. For example, the protocol might target emitting 7.5% of total supply in the first year, 6% in the second, and so on, asymptotically approaching the full allocation. This approach mirrors Bitcoin's philosophy of front-loading rewards to early adopters (who take on more risk and contribute when network effects are weakest), but with a smoother curve to continuously incentivize new contributors even as the project matures.
- **Epoch-Based Distribution:** Emissions could be distributed in discrete epochs (e.g., weekly or monthly). In each epoch, a fixed number of **TRUST** (determined by the schedule) is released to various reward pools.
- **Allocation by Activity:** The newly emitted **TRUST** each epoch is split among different activities according to governance-set weights. Typical buckets include:
 - *Staking Rewards:* to users staking TRUST on identities/claims (essentially content curation rewards).
 - *Bonding Rewards:* to users bonding TRUST in the TRUST bonding contract to receive voting rights and secure network nodes.
 - *Liquidity Rewards:* to users providing liquidity in **TRUST** trading pools (if the token is to be liquid on exchanges).
- **Long-term Cap:** Rather than having a hard cap of maximum tokens, Intuition allows a small perpetual inflation (e.g., less than 1% per year after year 10) to continuously reward participation, subject to governance.

The emission formula provided earlier, $E(t) = E(0)(1 - \rho)^t$, is one way to model the decay. Another way is piecewise: e.g., "X **TRUST** per block for first N blocks, then reduced by Y%." For concreteness, assume Year 1 aims to emit 50 million **TRUST**, Year 2 aims 40 million,

Year 3: 32 million, Year 4: 25.6 million, etc. This geometric reduction ($\rho = 0.20$ as in the example) means each year 80% of the previous year's emission is released. Over infinite time, it sums to a finite limit (here ~250 million).

The exact numbers will be decided by governance or initially set as in the token distribution table. What's important is **how emissions reward long-term alignment**:

The exact numbers will be decided by governance or initially set as in the token distribution table. What's important is **how emissions reward long-term alignment**:

- Early participants enjoy higher absolute rewards, but they are encouraged to **continue holding or staking** those tokens to maximize future gains.
- If a participant simply sells their **TRUST** immediately, they lose out on the compounding influence and rewards.
- Those who remain engaged (re-staking their rewards, participating in governance, etc.) effectively gain a larger share of the diminishing emission pie over time.

This model creates an implicit incentive for **loyalty**: as emissions drop, the value of being an active participant potentially increases (since fewer new tokens are shared among presumably more users, each token could be scarcer and more valuable, and the relative reward for staying might be the ongoing platform fees or other benefits rather than high inflation).

Bonding Mechanisms

Bonding in the context of Intuition refers to mechanisms where participants can lock up TRUST in exchange for rewards that typically vest over time. Bonding is about **commitment**: it asks participants to take on some temporary illiquidity or risk, rewarding them for the patience and long-term support.

Time-Locked Staking (Lockup/Bonding of TRUST): Holders of **TRUST** can choose to lock their tokens in the protocol for a fixed duration (say 3 months, 6 months, 1 year, up to a maximum of 4 years). In return for locking (or "bonding") their tokens, they receive boosted benefits – for example, increased voting power and a higher share of fee revenues or emissions. This is akin to the *vote-escrow* model popularized by protocols like Curve (veCRV) or Balancer (veBAL), where a longer lock yields more influence. In Intuition, we might call the locked token **veTRUST** (for "vested **TRUST**" or "vote-escrowed **TRUST**"). The boost could be linear with time: locking for the max period gives a 100% boost (2× rewards, or full voting

weight), while shorter periods give proportionally less.

Time-Locked Bonding (veTRUST)

When a user bonds (locks) their **TRUST**, they receive in exchange a non-transferrable *bonded stake representation*. Let's denote by $B(u, d)$ the *bonded weight* a user u gets for locking 1 **TRUST** for duration d . We can define, for example:

$$B(u, d) = 1 \times f(d),$$

where $f(d)$ is a scaling function with $f(0) = 0$ (no lock, no bonded weight beyond the base token itself) and $f(T_{\max}) = X$ (locking for the maximum duration yields X times the base weight). A common choice is linear scaling: $f(d) = 1 + \frac{X-1}{T_{\max}}d$ for $0 \leq d \leq T_{\max}$. If we set $X = 4$ (meaning a 4-year lock yields 4× weight, similar to Curve's 4-year = 2.5× model or others), then a 2-year lock (half of max) would yield $1 + \frac{3}{4} * 2 = 2.5$ times weight, etc.

This bonded weight has two main effects:

- **Governance Power:** A user's voting power in governance is based on bonded weight. This prevents a scenario where someone could buy tokens, vote immediately, and sell – they must lock tokens to have meaningful influence, aligning them with the long-term fate of the protocol. Thus, governance is controlled by the most committed holders.
- **Boosted Rewards:** Intuition allocates a portion of its protocol emissions and platform fees to bonded **TRUST** holders. The distribution of that share can be weighted by bonded weight. For example, if Alice locks 100 **TRUST** for max duration (getting 400 weight units) and Bob locks 100 **TRUST** for half duration (say 200 weight units), Alice will get twice the share of the fee distribution as Bob, reflecting her stronger commitment.

To illustrate, consider *Table 3*, which shows an example bonding reward multiplier for different lock durations (with $X = 2$ for simplicity in this table, meaning max lock doubles the benefits):

Bonding Period (Lock Duration)	Reward Multiplier	Voting Power Multiplier
No lock (0 weeks)	1.0× (baseline)	1× (can vote, but minimal influence)

Bonding Period (Lock Duration)	Reward Multiplier	Voting Power Multiplier
13 weeks (3 months)	1.2×	1.2×
26 weeks (6 months)	1.4×	1.4×
52 weeks (1 year)	1.7×	1.7×
104 weeks (2 years)	2.0×	2.0×
208 weeks (4 years, max)	2.0× (cap)	2.0× (cap)

Table 3: Example Bonding/Lockup Durations and Corresponding Multipliers.

In this hypothetical scheme, locking **TRUST**\$ for any period up to 2 years yields a proportional boost up to 2×, and any lock longer than 2 years is capped at 2× (just as an example; actual max lock might be 4 years with higher boost). The principle is that **longer locks = greater share of influence and rewards**. Users can choose a lock period that matches their confidence and liquidity needs; those who lock longer are taking more risk (their tokens are illiquid and exposed to protocol outcomes), so they are rewarded more.

The **mathematical notation** for reward distribution to bonded stakers could be as follows: suppose the protocol sets aside R tokens per epoch for bonded stakers as a whole (this could be from emissions or fee revenue). If user u has w_u bonded weight (after applying $f(d)$ to their locked amount), and the total bonded weight of all users is W_{total} , then u 's reward in that epoch is

$$R_u = R \times \frac{w_u}{W_{\text{total}}}.$$

This ensures a fair pro-rata distribution by weighted commitment.

System Utilization

To unlock the full bonding reward pool in each epoch, the **Intuition protocol** must demonstrate sufficient overall usage of its knowledge graph. *System Utilization* is a metric that measures this ecosystem-wide activity. An **epoch** (for example, a 2-week period, configurable via governance) is the time frame over which both usage and rewards are evaluated. For 100% System Utilization in an epoch, the network needs to see a net increase of **Y TRUST**

tokens deposited into the shared knowledge graph by all participants. This threshold **Y** is dynamically calculated based on the total TRUST emissions from the previous epoch – essentially requiring that a portion of last epoch’s rewards be put to productive use in the system.

- **Full (100%) Utilization:** If community contributions meet or exceed **Y** within the epoch, the system is considered fully utilized. In this case, the protocol releases the **maximum bonding rewards** for that epoch, as all newly emitted tokens are effectively being funneled back into the ecosystem’s growth.
- **Partial Utilization:** If the total new deposits fall short of **Y**, System Utilization will be proportionally lower (e.g. 50% of the target deposits achieved results in 50% System Utilization). The bonding reward emissions for that epoch are **scaled down** accordingly. For instance, if only half of the expected TRUST got deposited into the knowledge graph, only around half of the maximum rewards would be distributed to bonders (subject to the minimum floor discussed below). This creates an automatic brake on inflation whenever ecosystem usage lags.
- **Under-Utilization Floor:** Even in the extreme case of **0% System Utilization** (no meaningful increase in TRUST deposits that epoch), bonders are still guaranteed a baseline reward – for example, **25% of the normal reward**. This minimum ensures bonders aren’t completely penalized if the platform has a slow period, but such a scenario strongly signals underutilization. The protocol would respond by significantly **reducing emissions** until utilization improves, preventing excess token issuance from flooding the market with no accompanying growth in the knowledge graph.

Overall, the System Utilization requirement ties the platform’s inflation rate directly to actual usage. If usage of the Intuition knowledge graph doesn’t keep up with token emissions, new issuance contracts automatically. This **prevents inflationary pressure** from building up without corresponding ecosystem activity. It aligns the network’s growth trajectory with real adoption: **the more the system is used, the more rewards can be paid out**, promoting a positive feedback loop between token incentives and platform utilization. Conversely, if usage stagnates, the rewards throttle down to protect the token economy’s health and encourage renewed engagement, fostering **sustainable growth** over time.

Personal Utilization

While System Utilization looks at the big picture, *Personal Utilization* focuses on each

individual bonder's contributions. To receive their **full share of bonding rewards**, a user must remain an active participant in the Intuition ecosystem during the epoch. In practice, this means two things for every bonder in that period:

1. **Active Engagement:** The user must be active in the knowledge graph throughout the epoch (e.g. over the 2-week cycle). "Active" typically means making meaningful contributions to Intuition's token-curated knowledge graph – for example, adding new data (Atoms/Triples) or endorsing and curating existing data. This condition ensures the user isn't simply passively bonding their tokens; they are contributing to the platform's content and curation continuously.
2. **Recycling Rewards via Deposit (the X Requirement):** The user needs to deposit a certain amount of TRUST (call it **X**) into the knowledge graph during the epoch. Importantly, **X is not a fixed number**, but is dynamically determined based on the user's own bonding activity – specifically as a function of how much TRUST they have bonded and how many TRUST tokens they claimed as rewards in the last epoch. In essence, the more a user earned from bonding rewards previously, the more they are expected to **re-invest** back into the network's knowledge graph in the current epoch. This might involve staking those TRUST tokens to create or validate content in the graph (thereby turning their rewards into useful data contributions). For example, if a user bonded a large amount and earned 100 TRUST in rewards last epoch, the protocol might require them to deposit, say, 20 TRUST into new knowledge graph entries in the current epoch to maintain 100% Personal Utilization. A user with smaller bonded stakes and rewards would correspondingly have a smaller X to meet.

If the bonder satisfies both criteria above in the epoch, they achieve **100% Personal Utilization** and remain eligible for **their full allotment of bonding rewards** (assuming the System Utilization condition is also fully met). If they fall short – for instance, perhaps they only deposited half of the required X back into the graph, or they weren't active in creating any new data – then their Personal Utilization might be, say, 50%. In that case, the rewards they can claim for bonding will be **proportionally reduced** (again with a safety net floor in place). In other words, a user's reward is **directly tied to their personal contribution**: a bonder who only fulfills half of their expected utilization will receive roughly half of the potential reward amount (the exact scaling can be fine-tuned by the protocol, but the principle is that it's linear or at least monotonic with their utilization percentage). Even a completely inactive bonder (0% Personal Utilization) isn't left empty-handed due to the minimum reward floor (e.g. they'd still get ~25% of their base reward), but they would be foregoing the majority

of what they could have earned by participating fully.

The Personal Utilization mechanism creates a powerful incentive for every token holder who stakes (bonds) their TRUST to also **use the Intuition system actively**. It closes the loop between staking and usage: rather than just collecting passive inflation, bonders are encouraged to put those newly minted tokens to work by curating knowledge. This way, **every bonder becomes a user**, and the growth in token supply directly fuels growth in the knowledge graph's breadth and depth. The requirement to recycle a portion of one's rewards back into the platform mitigates sell pressure and idle accumulation, because users have a reason to *redeploy* tokens into Intuition's ecosystem (earning them further reputation or future rewards from the content they add). In sum, Personal Utilization aligns individual incentives with the network's health by making active contribution a prerequisite for maximum reward – rewarding those who help build the ecosystem and discouraging pure speculators who might otherwise stake and disengage.

Example Formulae (Illustrative)

Below is one simple parameterisation that makes the *Personal Utilisation* and *System Utilisation* scores fully explicit. Governance can change the coefficients (θ , ϕ , m_{floor}) at any time, but the logic stays the same.

Symbol	Meaning (epoch t)
E_{t-1}	Total TRUST emitted in the previous epoch
$R_{i,t-1}$	Reward actually claimed by user i in the previous epoch
B_i	Amount of TRUST that user i has bonded (principal)
$D_{i,t}$	TRUST deposited into the graph by user i this epoch
D_t^{sys}	Aggregate TRUST deposits by all users this epoch

1. Personal Utilisation

1. Required deposit threshold

$$X_{i,t} = \theta R_{i,t-1}$$

θ is a governance-set recycling coefficient (e.g. $\theta = 0.20 \rightarrow$ recycle 20 % of last-epoch rewards).

2. Personal score

$$U_{i,t}^{\text{personal}} = \begin{cases} \min\left(1, \frac{D_{i,t}}{X_{i,t}}\right) & \text{if the user made } \geq 1 \text{ graph interaction in epoch } t, \\ 0 & \text{otherwise.} \end{cases}$$

A fully active user who meets or exceeds the $X_{i,t}$ deposit target attains $U_{i,t}^{\text{personal}} = 1$.

2. System Utilisation

1. Network-wide deposit target

$$Y_t = \phi E_{t-1}$$

ϕ is a utilisation coefficient (e.g. $\phi = 0.30 \rightarrow$ community should recycle 30 % of last-epoch emissions).

2. System score

$$U_t^{\text{system}} = \min\left(1, \frac{D_t^{\text{sys}}}{Y_t}\right)$$

The target is hit—and $U_t^{\text{system}} = 1$ —whenever aggregate deposits match or exceed Y_t .

3. Reward Multiplier per User

Let

- m_{floor} be the protocol's guaranteed minimum share (e.g. $0.25 = 25\%$).

Then each bonder's effective multiplier in epoch t is

$$m_{i,t} = m_{\text{floor}} + (1 - m_{\text{floor}}) U_t^{\text{system}} U_{i,t}^{\text{personal}}.$$

Examples (with $m_{\text{floor}} = 0.25$)

U_t^{system}	$U_{i,t}^{\text{personal}}$	Reward share $m_{i,t}$
1.00	1.00	1.00 × (full reward)
0.60	1.00	$0.25 + 0.75 \times 0.60 = 0.70 \times$
1.00	0.50	$0.25 + 0.75 \times 0.50 = 0.625 \times$
0.00	0.00	0.25 × (floor)

In summary:

- **System activity** governs how much of the global reward pool is released.
- **Personal activity** determines each bonder's slice of that pool.
- **Neither score can exceed 1**, so rewards never surpass the scheduled maximum, and the **25 % floor** guarantees that even in a completely idle epoch bonders still receive something—though only a quarter of what they could have earned by hitting full utilisation.

Node Operator Delegation

TRUST bonders also optionally choose to delegate their bonded TRUST to node operators, to share in their risks and rewards. This functionality is outlined in detail below, in the **Node Operator Bonding and Slashing Requirements** section.

Aligning Incentives for Sustainable Growth

By combining System and Personal Utilization requirements, the bonding mechanism ensures that **to earn 100% of the possible rewards, two conditions must be met: the network as a whole is growing (System Utilization = 100%) and the individual participant is contributing (Personal Utilization = 100%)**. If either the user or the system under-utilizes the token (or both), the rewarded emissions automatically scale down. This dual-gauge approach has several important benefits for the Intuition protocol's tokenomics:

- **Prevents Unchecked Inflation:** Inflationary reward programs can undermine a token's value if rewards aren't tied to productive use. Here, any lack of utilization (individual or aggregate) curbs the issuance of new TRUST. This prevents excess tokens from hitting

the market unless there's actual value being created in tandem (new knowledge graph data, more usage of the platform). In effect, it **avoids inflationary pressure without corresponding ecosystem usage**.

- **Ensures Utilization of Rewards:** The model explicitly encourages and **incentivizes users to recycle rewards** into the system. Rather than viewing rewards as profit to cash out, users see them as fuel to further engage with the knowledge graph (which can create compounding benefits, as useful contributions can generate additional reputation or future rewards in Intuition). This means the token economy drives real activity: staking and contributing become part of the same feedback loop.
- **Aligns Individual and Collective Incentives:** Every participant's full earning potential is contingent not just on their own actions but also on the community's overall activity. This alignment motivates users to not only contribute personally but also **promote broader usage** – for instance, by onboarding others or improving the quality of shared data – so that the whole network meets its utilization goals. Bonders know that if the platform thrives, everyone's rewards grow, creating a sense of shared responsibility for Intuition's success.
- **Promotes Sustainable Growth:** During times of high engagement and expansion of the knowledge graph, the protocol generously rewards contributors (distributing the maximum scheduled emissions). During quieter periods, it slows down token distribution, giving the ecosystem time to catch up. This adaptive issuance model helps the project grow organically and **sustainably**, minimizing boom-bust dynamics. It focuses on long-term health over short-term reward chasing.

The **System & Personal Utilization** mechanisms embedded in Intuition's bonding process ensure a healthy, use-driven economy. They act as safeguards that align token inflation with actual network value creation. Even though a minimum reward (e.g. 25% of the full rate) is always provided to keep baseline incentives in place, only those participants who actively engage with Intuition and contribute to the knowledge graph – and do so in an era of growing overall usage – will unlock the **full 100% of bonding rewards**. This approach keeps **incentives properly aligned**: token holders are motivated to become true stakeholders in the ecosystem's growth, and the platform's expansion goes hand-in-hand with its token distribution. Ultimately, tying rewards to utilization helps cultivate a vibrant, **self-reinforcing cycle** of participation, data creation, and value accrual, driving Intuition toward long-term prosperity.

Decentralized Indexing & Query Service

In Intuition, node operators play a pivotal role in decentralizing data availability and query processing. Each operator serves as a decentralized indexer, maintaining a portion of the Intuition knowledge graph so that it can be efficiently queried by users and applications. This is necessary due to the fact that the Intuition knowledge graph is comprised of both on-chain and off-chain state—the nodes, edges, and weights of the knowledge graph are stored on-chain, and the heavy metadata referenced is stored off-chain. Intuition nodes compose together the on-chain and off-chain state into a single, traversable and queryable graph. The allocation of indexing responsibilities follows a hybrid approach that balances network resilience with operator specialization.

Redundancy and Load Assignment

Intuition's indexing network **randomly assigns knowledge graph segments** to multiple nodes to ensure both redundancy and balanced load. Each data segment is replicated across a subset of nodes chosen via uniform random sampling. For example, if each segment is assigned to r out of N total nodes (uniformly at random), then any given node has probability $\frac{r}{N}$ of indexing a particular segment. This means the **expected number of nodes** covering the same data segment is $E[\text{nodes per segment}] = r$, providing a high degree of redundancy. Moreover, any two distinct nodes have a probability

$$P(\text{both index same segment}) = \frac{\binom{r}{2}}{\binom{N}{2}} = \frac{r(r-1)}{N(N-1)},$$

of overlapping on a given segment. This overlap ensures that even if one indexer goes offline, others can serve the same data, maintaining query availability.

This **randomized assignment** also spreads indexing responsibility evenly. With S total segments in the knowledge graph, and each segment indexed by r nodes, the total index assignments in the network is $r \cdot S$. Each of the N nodes will index roughly $\frac{rS}{N}$ segments on average, smoothing out load. In fact, the load distribution can be modeled as a binomial process – the number of segments a node indexes follows approximately a Binomial $(S, r/N)$ distribution – which for large S concentrates around the mean $\frac{rS}{N}$. This probabilistic load balancing prevents any single node from becoming a bottleneck and helps the network scale as more data and nodes are added.

Uptime Enforcement

To guarantee reliability, Intuition enforces a strict **uptime requirement** for indexing nodes. Each epoch (a fixed time period), a node's uptime u is measured as the fraction of time it was online and responsive (e.g. $u = \frac{T_{\text{online}}}{T_{\text{epoch}}}$). Nodes are expected to maintain u above a minimum threshold θ (for instance, $\theta = 0.95$ meaning 95% uptime). If a node's uptime in an epoch falls below this threshold ($u < \theta$), a **partial slashing penalty** is applied to discourage poor availability. We define a slashing function $S(u)$ that deducts a portion of the node's stake or rewards proportional to the shortfall in uptime. For example, a simple linear penalty model can be:

$$S(u) = \begin{cases} 0, & u \geq \theta, \\ \lambda(\theta - u), & u < \theta, \end{cases}$$

where λ is a scaling constant that determines the severity of the penalty. Under this scheme, a node that barely misses the target might lose a small fraction, whereas a node with significantly low uptime would be penalized more heavily. This incentivizes operators to maintain high availability; consistent downtime directly translates into financial loss. By adjusting parameters like θ and λ , the network can tune how strict the uptime enforcement is, ensuring that **indexers remain online** and queries are served reliably.

Reward Distribution

Intuition's economic model uses a two-pronged reward system to incentivize honest and efficient behavior: **protocol emission rewards** and **query fee rewards**. All rewards are distributed in proportion to each node's contributions and performance, aligning the indexers' incentives with network health.

Emission Rewards

The protocol allocates a certain number of native tokens each epoch (the emission pool) as rewards for indexers. These **emission rewards** are distributed proportionally based on several factors, including the amount of data a node has indexed (e.g. number of knowledge graph claims indexed), the volume of user queries it has processed, and its reliability metrics (such as recent uptime and correctness of query responses). We can formalize the allocation as follows: let C_i , Q_i , and ρ_i denote node i 's metrics for indexed content, queries processed,

and reliability score, respectively, over the epoch. Assign weighting coefficients α , β , and γ to reflect the importance of each factor. Then define a performance score for node i as

$$P_i = \alpha C_i + \beta Q_i + \gamma \rho_i ,$$

aggregating its contributions. If R_{epoch} is the total tokens available for emission rewards in that epoch, node i 's reward share is:

$$R_i = R_{\text{epoch}} \frac{P_i}{\sum_j P_j} ,$$

where the sum in the denominator runs over all active indexer nodes j . In other words, each node earns a fraction of the reward pool equal to the fraction of the total performance score it contributed. This formula ensures **fair reward distribution** – a node that indexes more data and handles more queries (and does so reliably) will naturally receive a larger portion of the emissions. High reliability (ρ_i close to 1) boosts a node's score, so flaky nodes with low uptime or accuracy earn substantially less. Over time, this proportional system encourages nodes to maximize their useful work and maintain integrity to earn more of the emission rewards.

Query Fees

In addition to emissions, indexers earn **query fees** – micro-payments from users or dApps for serving queries. Intuition's service aggregates these fees into a pool and then distributes them among the nodes based on the actual query traffic each node handled, adjusted by the quality of service they provided. Each node i accrues a certain number of queries Q_i served in the epoch, but not all queries are equal: low-latency, correct responses are valued more. We introduce a performance weighting factor ϕ_i for query handling, which captures the node's **quality of service** (e.g. combining average latency and response accuracy into a single score). For instance, one simple model could be $\phi_i = \frac{A_i}{L_i}$, where A_i is the fraction of queries answered correctly (accuracy) and L_i is the node's average response latency – this gives higher weight to nodes that are more accurate and faster. Using such a factor, the fee distribution can be computed similarly to emissions: if F is the total fee pool collected from users in an epoch, then node i 's share of query fees is

$$R_i^{(fee)} = F \frac{Q_i \phi_i}{\sum_j Q_j \phi_j} ,$$

summing over all indexers j . This formula means fees are allocated in proportion to **traffic served, weighted by performance**. A node that answers many queries will earn more, but only if it maintains good performance (fast and correct responses). If two nodes serve the same number of queries, the one with lower latency or higher accuracy (higher ϕ) will receive a bigger share of fees. This mechanism incentivizes nodes not just to handle volume, but to continually improve response quality, benefiting end users with faster and more reliable queries.

Query Routing Optimization

When a user's query is submitted to Intuition's network, the protocol must decide **which node** will execute the query. Rather than picking an indexer at random, Intuition employs a weighted routing strategy that favors nodes with a proven track record. Each node i is assigned a **routing weight** w_i based on its historical performance (e.g. recent uptime u_i and response accuracy a_i). For example, the network might compute w_i as the product $w_i = u_i \times a_i$, treating both uptime (as a fraction of the last epoch, for instance) and accuracy (fraction of queries answered correctly) as values between 0 and 1. A well-performing node with $u_i \approx 1$ and $a_i \approx 1$ gets a weight close to 1, whereas a node with 80% uptime and 75% accuracy would have $w \approx 0.8 \times 0.75 = 0.60$, relatively lower. Once weights are determined, **query routing** is done probabilistically in proportion to these weights. The probability that a new incoming query is assigned to node i can be modeled as:

$$P(\text{node } i \text{ serves next query}) = \frac{w_i}{\sum_j w_j},$$

where the sum is over all candidate indexer nodes j currently available to serve. This weighted random selection ensures that nodes with higher performance histories are more likely to receive queries, thereby improving overall network latency and correctness. Importantly, it's still probabilistic – less performant nodes aren't completely excluded, which allows new or improving indexers to gradually earn traffic as they build up reputation. Over time, this **performance-weighted routing** optimizes query resolution by directing queries to the most reliable nodes (increasing the chance of fast and accurate responses) while still load-balancing across the network. It creates a positive feedback loop: nodes are motivated to keep their uptime high and error rates low in order to receive more queries (and hence more fees), and users experience a faster, more dependable query service as a result.

Node Operator Bonding and Slashing Requirements

Bonding Requirement: Every node operator on the Intuition network is required to **bond a fixed amount of TRUST tokens** as collateral, denoted as the `BOND_REQUIREMENT`. This bonded stake serves as the operator’s “skin in the game,” aligning their incentives with the network’s integrity. A node’s bond can be satisfied entirely by the operator’s own self-staked TRUST or by a combination of self-stake and tokens **delegated** (bonded) from other holders (referred to as *bonders* or delegators). In either case, the total bonded collateral for a node must **always remain at or above** the `BOND_REQUIREMENT`. If a node’s total bond falls below this threshold – for example, due to slashing penalties or withdrawal of delegated tokens – the node immediately becomes ineligible to operate. Such a node is **removed from the active indexing set**, meaning it can no longer participate in query processing or earn network rewards until the bond is restored to at least the minimum requirement.

Slashing Conditions and Tiered Penalties: To encourage reliable service, Intuition implements a **slashing mechanism** that penalizes node operators for misbehavior or underperformance. Slashing events cause a **deduction of a fraction of the node’s bonded TRUST**, directly reducing the operator’s staked collateral (and thus that of its delegators). Slashing is **tiered and proportional** to the severity and frequency of violations: minor or infrequent lapses incur small penalties, whereas severe or repeated failures lead to larger slashes. We define a general slashing penalty function: $S_{\text{violation}}(\text{type}, \text{severity})$, which computes the slashed fraction of bonded tokens based on the violation type and its severity. Different types of misbehavior – such as prolonged downtime, incorrect query responses, or lack of responsiveness – map to different penalty tiers within this function. Below we outline the key slashing conditions and their penalty formulations:

- **Uptime Violation:** Each node must maintain a minimum uptime θ (fraction of time online). If a node’s actual uptime u drops below this threshold ($u < \theta$), the network imposes a slashing penalty proportional to the shortfall. We preliminarily maintain the existing uptime-based slashing formula for examples sake, to rudimentarily demonstrate the mechanic:

$$S_u = \lambda (\theta - u) \quad \text{for } u < \theta,$$

where $0 < \lambda \leq 1$ is a coefficient controlling the slashing severity. For example, if $\theta = 0.95$ (95% required uptime) and an operator only achieved $u = 0.94$, then $\theta - u = 0.01$ (1% shortfall). A chosen λ (say $\lambda = 1$ for a full-proportional penalty) would result in

$S_u = 0.05$, meaning 5% of the node's bonded tokens are slashed. Larger deviations from the uptime target or repeated downtime episodes should trigger higher λ or additional penalties, reflecting the greater severity.

- **Incorrect Response (Accuracy) Violation:** Nodes are expected to serve **correct data** consistently. We define each node's response accuracy A_i as the fraction of data it returns from queries correctly (as determined by verification or consensus), and we set a minimum accuracy threshold ϕ (e.g. 99%). If $A_i < \phi$ over a given evaluation period, the node faces a slashing penalty. The slashed fraction can be computed in proportion to the accuracy shortfall, for instance:

$$S_A = \alpha (\phi - A_i) \quad \text{for } A_i < \phi,$$

where α is a penalty coefficient for accuracy-related violations. This means a node that serves significantly incorrect data (failing to meet the accuracy threshold) will lose a portion of its bonded tokens commensurate with how far its accuracy fell below ϕ . The network may also tier this penalty by severity – e.g. a slight dip below ϕ yields a small slash, whereas egregiously incorrect data or deliberate misinformation (very low A_i) results in a much larger slash. This incentivizes operators to cross-verify data and maintain correctness, as serving wrong information has direct financial consequences.

Each of these slashing functions S_u and S_A , deducts the calculated fraction of the node's **total bonded collateral**. All slashing events directly reduce the node's bonded amount in real time. If, after any slashing event, the node's remaining bonded stake falls below the required `BOND_REQUIREMENT`, the protocol **automatically ejects the node from the active indexing set**. This ejection is immediate and non-discretionary – it serves as a safety mechanism to ensure that only well-collateralized and compliant nodes continue to serve the network. An ejected node must replenish its bond (back to at least the minimum requirement) and possibly face a cooldown or re-application process before it can resume operation.

Shared Slashing Risk for Delegators (Bonders): One important aspect of Intuition's design is that **both node operators and their bonders share the slashing risk** in proportion to their contributions to the bond. When a slashing event occurs for a given node, the penalty (loss of bonded tokens) is distributed across all staked funds for that node. For example, if a node with a total bonded stake B (including operator's own stake and all delegated stakes) incurs a slashing of 10% ($S_{\text{violation}} = 0.1$), then **every participant's bonded amount is reduced by 10%**. A delegator who bonded 100 TRUST to that node would lose 10 TRUST, the same 10% fraction as the operator loses on their self-bonded stake. In this way, *bonders incur slashing*

proportional to their share of the node's bonded collateral. This shared-risk model ensures that delegators are not passive bystanders but are economically incentivized to delegate to reliable, well-performing nodes. Before bonding their TRUST to a node, token holders will assess the operator's track record for uptime, accuracy, and responsiveness – because poor performance could put their delegated stake at risk. Bonders are thus aligned with operators in seeking to maximize performance and avoid penalties.

Incentive Alignment: The bonding and slashing requirements create a robust incentive structure: node operators must **lock a substantial amount of value** (the `BOND_REQUIREMENT` in TRUST tokens) which they stand to lose if they fail to meet the network's performance standards. Delegators, by contributing to a node's bond, **share in both the rewards and the risks**. If the node performs well, both operator and delegators earn query fees and network rewards; if the node performs poorly, both sides are penalized. This mutual stake in the node's success encourages vigilant maintenance, prompt software updates, accurate indexing, and reliable uptime. In summary, bonding and slashing mechanisms work in tandem to **enforce accountability**: they deter malicious or negligent behavior by making it financially costly, and they reward diligent operators who consistently meet or exceed the network's reliability thresholds. Ultimately, this strengthens the Intuition protocol's security and trustworthiness by ensuring that only highly committed and performant nodes thrive in the network ecosystem.

By introducing decentralized node operators into the utility mechanisms of the network, Intuition significantly decentralizes its infrastructure layer. Instead of relying on any single entity or core team to host and serve the entire knowledge graph (which consists of both on-chain and off-chain state), the data is distributed across many independent operators around the world. This decentralization hardens the system against censorship and downtime: no centralized party can shut down access to information, and the failure of any single node (or even many nodes) does not compromise the availability of the knowledge graph. It also allows Intuition to scale organically; as demand for knowledge graph queries grows, more node operators can join and share the load, each motivated by the opportunity to earn rewards. Importantly, these incentives tie together the interests of data curators, node operators, and token holders in a harmonious way. Curators who stake TRUST to add high-quality data see that data widely available and reliably served by operators, increasing its usage and the rewards it generates. Node operators earn more by hosting the data that users find most valuable, which often corresponds to data curated with high TRUST stakes and community endorsement. In turn, as query fees and utility of the network increase, demand for the TRUST token (for staking, fees, and rewards) should rise, benefiting all token holders. In

summary, the node operator role strengthens Intuition's utility layer by ensuring data availability and integrity at scale, while further aligning incentives across the ecosystem – every actor, from curator to operator to end-user, is rewarded for making the knowledge graph more reliable, accessible, and valuable.

Governance Framework

The Intuition platform is governed as a decentralized protocol, with **TRUST token holders** at the center of decision-making. Governance in Intuition is designed to be **progressively decentralized**, with more and more control shifting to the community of **TRUST** holders over time. The governance framework addresses two main domains: **protocol parameters** (economic and functional settings that define how Intuition operates) and **token emissions allocation** (deciding how newly minted **TRUST** is distributed).

Governance Model and Process

Intuition employs a **DAO (Decentralized Autonomous Organization)** model for governance. Every **TRUST** token represents a voice in the DAO, and key decisions are made via proposals and votes. The governance process typically follows these stages:

1. **Proposal Creation:** Community members can draft proposals to change protocol settings, introduce new features, or allocate treasury funds. To prevent frivolous proposals, a minimum threshold of **TRUST** (or *delegated TRUST*) is required to submit a proposal for a formal vote. This threshold ensures only those with sufficient stake in the system can move an issue to the ballot. (For instance, one might need 1% of total supply backing the proposal to proceed to voting, though this exact number can itself be tuned by governance.)
2. **Discussion and Off-Chain Signaling:** Proposals are discussed on governance forums or off-chain platforms. Here, token holders (and even non-holders) debate the merits. Off-chain signaling tools (like Snapshot votes or social polls) might gauge sentiment before an on-chain vote. This phase leverages community insight to refine proposals. It is common for proposals to be iterated based on feedback before final submission.
3. **On-Chain Voting:** Using **TRUST** tokens (often locked or delegated to voting contracts), holders vote on proposals. Each token typically equals one vote (1 **TRUST** = 1 vote) by

default. Intuition may also implement **quadratic voting** or **boosted voting** (vote-escrow) mechanisms to balance influence – for example, requiring users to lock their tokens for a period (bonding, see next section) to gain full voting power. If a proposal achieves a majority (and meets any quorum requirement, e.g. at least 10% of total supply participated), and no vetos or safety modules object, the proposal is considered passed.

- 4. **Execution:** Intuition is built on smart contracts, so many governance decisions can be executed on-chain automatically. Passed proposals trigger changes such as adjusting a parameter in a contract, or transferring funds from the treasury, via a governance executor contract. In early stages, a multi-signature of trusted community members or team may act as a safeguard, but the end goal is autonomous execution.

Delegated Governance: Recognizing not all token holders will actively vote, Intuition's framework supports delegation. **TRUST** holders can delegate their voting power to representatives or experts (often called "delegates") who vote on their behalf. This improves governance participation and decision quality, as active delegates can accumulate voting weight and develop expertise in protocol issues. Delegation is flexible; holders can change or revoke their delegate at any time. This approach draws inspiration from governance systems of protocols like Compound and Uniswap, aiming for a balance between direct decentralization and practical voter engagement.

Governance Scope: Parameters and Controls

TRUST token holders have the authority to govern a wide range of parameters and elements of the Intuition protocol. *Table 2* outlines some of the key governance rights and responsibilities conferred via **TRUST** ownership:

Governance Domain	TRUST Holder Rights
Economic Parameters	Adjust fee rates for platform actions (e.g., identity creation fee, staking fee or reward commission). Vote on staking reward parameters such as stake lockup periods, cooldown times, or slashing rules for malicious operations. Tune the emission rate ρ used in the token emission formula or even enact a change in total supply (e.g., introduce a burn or a new token mint for specific events) with broad consensus.

Governance Domain	TRUST Holder Rights
Emissions Allocation	Allocate weekly or monthly TRUST emissions to various incentive programs . For example, decide what percentage of new tokens go to content staking rewards versus liquidity mining or developer grants. This can be implemented via a gauge voting system where TRUST holders vote to weight different "gauges" (e.g., specific data categories or pools) and emissions are distributed proportionally. This ensures the community directs incentives to areas deemed most crucial for growth.
Treasury Management	Propose and vote on the use of the Intuition Treasury (which holds funds in TRUST and possibly ETH from fees). Examples: funding community proposals or bounties, granting tokens to bootstrap new integrations or research, financing continued protocol development, or initiating buyback-and-burn programs to support token value. All spending from the treasury is subject to approval by token holder vote, enforcing accountability and alignment with TRUST holders' interests.
Protocol Upgrades	Approve upgrades to Intuition's smart contracts or the introduction of new modules (for instance, a new bonding curve contract or a new version of the staking system). Typically, major upgrades would be proposed by core devs or a technical committee and then ratified by the community. TRUST holders ensure upgrades maintain decentralization and security principles.
Governance Processes	Meta-governance rights to modify how governance itself works. For example, change the voting quorum or proposal threshold, choose to implement quadratic voting or change delegation rules, or even initiate the formation of sub-DAOs or councils for specialized decisions. This is essentially the community's ability to refine its own decision-making structure as the ecosystem grows.
Emergency Powers	In extreme cases, TRUST holders (or a council elected by them) may exercise emergency powers, such as pausing certain protocol functions during an attack or quickly adjusting a parameter to protect the system. These are used sparingly and usually require a higher threshold (e.g.,

Governance Domain	TRUST Holder Rights
	supermajority) or rapid multisig action, with subsequent community review.

Table 2: Governance Rights of **TRUST** Token Holders.

Through these rights, **TRUST** holders collectively steer Intuition. The design ensures that those who are economically invested in the network's success have control over the system, aligning **control with stake**. Notably, governance of Intuition is an evolving process: initial conditions (like emission rates or fee percentages) may be initially set by initial deployer(s) of the contract(s), but they are ultimately **subject to change by community vote**, allowing the system to adapt over time.

Emissions Voting and Parameter Tuning

Two governance activities warrant special emphasis given their importance to token economics: **emissions voting** and **parameter tuning**.

Emissions Voting — “Gauge”-Style Allocation

Intuition borrows the **gauge voting** design pioneered by Curve Finance, adapting it to the needs of a knowledge-graph economy rather than a pure AMM. Voting power comes from **veTRUST** (vote-escrowed TRUST), so only long-term-committed holders control where new emissions flow.

Symbol	Meaning (epoch t)
E_t	Total TRUST scheduled to be emitted in epoch t (after Utilization scaling)
$\text{veTRUST}_{i,t}$	Vote-escrow balance of user i at start of epoch t
$f_{i,g,t}$	Fraction of user i 's votes sent to “gauge” (pool/category) g this epoch, with $\sum_g f_{i,g,t} \leq 1$
$W_{g,t}$	Raw weight of gauge g after tallying votes

Symbol	Meaning (epoch t)
$\rho_{g,t}$	<i>Relative weight (share of emissions) for gauge g</i>
$E_{g,t}$	Emissions allocated to gauge g this epoch

1. Per-user vote casting

Each user splits their voting power across any number of gauges:

$$0 \leq f_{i,g,t} \leq 1, \quad \sum_g f_{i,g,t} \leq 1.$$

A holder who wants all new emissions to reward *claim verifiers*, for example, would set $f_{i,\text{verifiers},t} = 1$.

2. Gauge weight aggregation

Aggregate raw weight is simply the sum of veTRUST-weighted votes:

$$W_{g,t} = \sum_i (\text{veTRUST}_{i,t} f_{i,g,t}).$$

Total voting weight that epoch is

$$W_t^{\text{tot}} = \sum_g W_{g,t} = \sum_i \text{veTRUST}_{i,t}.$$

3. Relative weight and emission split

$$\rho_{g,t} = \frac{W_{g,t}}{W_t^{\text{tot}}}, \quad E_{g,t} = E_t \times \rho_{g,t}.$$

Thus, if a gauge attracts 12 % of the vote weight, it receives 12 % of that epoch's TRUST emissions.

4. veTRUST “decay” (Curve-style)

A user’s veTRUST balance decays linearly to zero over their lockup period:

$$\text{veTRUST}_{i,t} = B_i \times \frac{d_i - t}{d_i - t_i^{\text{lock}}} \quad (0 \leq t \leq d_i),$$

where

- B_i = bonded principal,
- t_i^{lock} = lock start,
- d_i = lock expiry.

Longer locks therefore carry more enduring voting power, aligning votes with long-term interests.

5. Example

Suppose $E_t = 100,000$ TRUST for the coming epoch and three gauges exist as determined by the canonical Intuition ontology (also voted on through governance, which represent different regions of the knoweldge graph:

Gauge	Raw weight $W_{g,t}$	Relative $\rho_{g,t}$	Emission $E_{g,t}$
Web3-related data	2.0 M	40 %	40 000 TRUST
Consumer-products related data	1.5 M	30 %	30 000 TRUST
Science-related data	1.5 M	30 %	30 000 TRUST

Token holders can redirect future flows every epoch; if the network later decides data verification needs more juice, voters can shift weight accordingly.

6. Why gauge voting matters

- **Fine-grained capital steering:** Emissions continuously migrate to the activities the community considers most value-accretive.
 - **Self-correcting incentives:** As a category matures (e.g.\ identity bootstrapping is “good enough”), voters naturally shift weight to the next bottleneck.
 - **Bribe resistance via long-term locks:** Because veTRUST decays with time-remaining, short-term mercenaries have limited sway.
-

Parameter Tuning (Quick Recap)

Separate proposals let veTRUST holders update protocol constants. A generic on-chain vote sets a new parameter p :

$$p_{t+1} = p_t + \frac{\Delta p}{W_t^{\text{tot}}} \sum_i (\text{veTRUST}_{i,t} s_i),$$

where $s_i \in \{-1, 0, +1\}$ encodes each wallet’s stance (for/against/abstain). Parameters govern, for example:

- θ, ϕ — recycle-target coefficients in Utilization math
- **Fee splits:** treasury vs.\ content creators
- **Bonding multipliers:** slope and caps of $f(d)$

Routine, data-driven tuning keeps Intuition’s economics responsive without resorting to any central administrator.

Borrowing the proven gauge-voting mechanics of Curve gives Intuition a **battle-tested, fully on-chain way to direct emissions** to where they do the most good, while parameter votes let the community fine-tune the engine as real-world data rolls in.

Parameter Tuning

Intuition's economic model involves several parameters – such as fees, reward rates, and

bonding multipliers – that may need adjustment as real-world data comes in. Rather than hard-code these forever, the governance allows updating them. For example, if the initial identity creation fee is 0.05 **TRUST** but this proves too high (discouraging contributions) or too low (not filtering spam), governance can vote to change it. Similarly, the split of fees that go to prior contributors versus the treasury could be tuned. Parameter tuning is guided by both data (KPIs like growth of contributions, or abuse metrics) and the collective judgment of the community on what will optimize network health. By making these parameters governable, Intuition ensures it can **evolve** from experimental stages to a stable, efficient state without requiring centralized intervention.

Inclusivity and Off-Chain Inputs

While **TRUST** token voting is the formal mechanism, Intuition's governance framework is designed to be **inclusive and informed by a broad set of stakeholders**. This means encouraging participation from not just large token holders but also smaller holders and even external experts (via delegation). Off-chain discussions and advisory committees (such as a research council or an economic advisory board) could be established to provide analysis on proposals. For example, changes to emission rates might be accompanied by simulation data from economists, and changes to technical parameters might get input from core developers. Ultimately, **TRUST** holders make the decision, but they do so with transparent information.

This inclusive approach is critical in a system as complex as a token-curated knowledge graph. The interplay of incentives can be subtle, and decisions should be made with care to avoid unintended consequences. The governance framework therefore emphasizes *transparency, education, and robust debate*. All governance actions and their rationales are recorded, and over time a knowledge base of governance decisions (a "governance graph" attached to the knowledge graph) will form, itself an asset for future decision-makers.

In conclusion, the governance framework powered by **TRUST** ensures that Intuition is not a static protocol, but a living system responsive to its community. By giving token holders the reins to adjust emissions and parameters, and by providing a clear process for proposals and voting, Intuition can harness the collective wisdom of its user base – the very premise of its knowledge graph – to guide the platform's evolution. Next, we examine how **TRUST**'s supply is emitted and how a bonding mechanism reinforces long-term participation, complementing the governance model described here.

Economic Incentives and Alignment

The ultimate goal of Intuition's token economic design is to align incentives among all participants – data creators, curators (stakers), consumers, and token holders – such that the network grows in a healthy way and trustworthy data is consistently produced and maintained. In this section, we narratively summarize how the **economic incentives** built into **TRUST** and the Intuition protocol work together to create a self-sustaining, virtuous cycle. We also highlight how these incentives handle potential conflicts or edge cases.

Aligning Creators and Curators

Data Creators populate the knowledge graph by adding new identities and claims. Their incentive is twofold:

- They might personally want the data on the network (for utility, e.g. a developer publishing data for use in their application, or a user adding a piece of information they care about to their profile).
- Economically, by being the first to add a valuable piece of data, they gain a stake in it and stand to earn rewards as that data gains traction. This is akin to being an early content creator on a platform that later explodes in popularity – Intuition ensures those early creators capture some of that upside in a measurable way (through fee shares and token emissions).

Curators (Stakers), who might also be creators or separate participants, look for the most promising or accurate information to back with their stake. They are essentially performing a due diligence or signal amplification role. Their incentives:

- Earn returns on their stake: Each claim or identity with fees flowing through it generates yield for stakers. If a curator correctly identifies a claim that will become important (maybe it's a crucial fact or a popular topic) and stakes on it early, they will accumulate a larger share of that item's vault and therefore a larger portion of all future fees/rewards from it.
- Influence network knowledge: Curators also have an intrinsic motivation (in a well-designed system) to see truthful and high-quality information prevail, especially if their reputation or **TRUST** holdings are at stake. There's a social incentive: being known as a

top curator (with lots of stake on correct info) could confer status in the community, perhaps even governance clout if that reputation translates into elected delegate roles.

The interplay: creators often act as the first curator of their own data (by staking the fee they paid into it). Subsequent curators validate or challenge that data by staking as well. If a piece of data is **good** (accurate, useful), curators will add positive stake, raising its profile and rewards. If it's **bad** (false or not useful), ideally no one stakes on it, or someone stakes against it, and it withers due to lack of economic support. Thus, both creators and curators are financially motivated to populate the graph with information that others will eventually find and use. Useless data is a wasted opportunity (fees spent with no return), and harmful or false data can even lead to slashing or social consensus penalties, which are disincentives to avoid.

Long-Term Holders and Governance Participants

Token Holders (especially those who bond their **TRUST** long-term) have an interest in the overall health and success of the Intuition platform:

- They benefit from any potential token value appreciation which may be tied to network usage (as more people use Intuition, more fees are generated, more tokens get locked up for staking/governance, etc., potentially driving demand for **TRUST**).
- They directly benefit from **value capture mechanisms**: For instance, if a portion of fees is reserved for **TRUST** stakers or if tokens are bought back and burned, the economic value flows to committed token holders.
- Through governance, holders can actively shape policies to increase the value of the network – e.g., adjust fees to optimal levels, authorize partnerships, or improve incentive alignment by tweaking parameters.

Crucially, long-term holders who are locked into governance (via bonding) are incentivized to encourage sustainable practices rather than short-term extraction. Because their tokens cannot be immediately sold, they would lose if the ecosystem collapses or trust in the data is broken (which would render **TRUST** worthless). Therefore, they will vote for decisions that improve data quality, attract more users, and ensure the longevity of the platform (for example, maintaining moderate inflation rather than hyperinflation, or enforcing prudent treasury spending).

Developers and Integrators: An important long-term stakeholder group is developers building on Intuition (using its data via APIs, integrating the knowledge graph into applications, publishing data). While they may not directly earn TRUST by default (outside of the TRUST earned through creating, curating, and bonding), governance may opt to allocate grants or emissions to incentivize third-party developers. This is part of emissions allocation governance – the community might decide to reward those who build useful tools that increase network usage (via the treasury or via specific programs). Developers are aligned with token holders in that they want a rich, reliable data platform to build on, which in turn brings more usage and possibly fees (some of which loop back to token holders). In essence, TRUST can be thought of as giving a voice and a share in the success of the network to not just curators, but anyone contributing to network effects.

Fee Structures and Value Flow

One of the core alignment mechanisms is the **distribution of fees** from actions on the platform. Whenever a user performs a paid action (like creating a claim or staking on an existing claim), that fee is split among stakeholders and the protocol. This ensures continuous redistribution of value:

- Prior contributors get rewarded (which encourages *early participation*).
- The protocol (and by extension **TRUST** holders via the treasury) captures some value for ongoing development and support.
- New contributors gain ownership (they effectively convert their fee into stake, giving them future upside).

Table 4 provides a conceptual breakdown of how fees from various actions might be distributed, showing the economic alignment between different parties:

Action & Fee	New Contributor (Actor)	Prior Contributors (Stakeholders)	Treasury/ Protocol
Create new Identity (fee e.g.	Fee converts to stake for creator (e.g.	A portion of fee distributed to TRUST treasury or burned (small, e.g. 0.002 TRUST) as network revenue, since for brand new data there are no prior	Protocol takes a small listing fee (e.g. 0.001 TRUST) into

Action & Fee	New Contributor (Actor)	Prior Contributors (Stakeholders)	Treasury/ Protocol
0.01 TRUST)	0.007 TRUST becomes creator's stake in that identity's vault) – giving the creator future claim on rewards.	contributors yet.	treasury for facilitation. Early identity creation might have minimal prior stakeholders, so most of fee becomes the creator's stake, effectively.
Endorse an Identity (stake on existing Identity, fee e.g. 0.005 TRUST)	Endorser's fee becomes their stake (e.g. 0.0035 TRUST to stake) in the identity's vault, increasing their share of it.	The identity's prior stakers receive a cut (e.g. 0.001 TRUST split among them proportional to their stake, represented as a share price increase), rewarding them for vetting that identity early.	Protocol takes the remainder (e.g. 0.0005 TRUST) as revenue.
Create a Claim (new Triple) (fee e.g. 0.02 TRUST)	Creator's fee partly stakes the Triple (say 0.01 TRUST stake for	The stakeholders in each underlying Atom referenced receive a reward. For instance, if the claim uses an existing predicate Atom or links two Atoms, those Atom's stakeholders might receive a portion (e.g. 0.005 TRUST) as	A protocol fee (e.g. 0.005 TRUST) goes to treasury.

Action & Fee	New Contributor (Actor)	Prior Contributors (Stakeholders)	Treasury/ Protocol
	creator in the claim's True vault).	the claim increases the utility of their identities. This cross-rewards the maintainers of connected knowledge.	
Stake on an existing Claim (fee e.g. 0.01 TRUST)	Staker's amount primarily goes to increase their stake (e.g. 0.008 TRUST to vault).	Prior stakers of that claim receive a reward (e.g. 0.0015 TRUST split) as a "dilution incentive" – by rewarding earlier stakers when new stake comes in, it mimics how early equity holders are rewarded when new investors join at a higher valuation. This incentivizes people to stake early on claims they think others will eventually also stake on. Stakeholders in each underlying Atom also receive a reward, as this claim stake signals relevancy of each underlying Atom.	Protocol fee (e.g. 0.0005 TRUST) to treasury.
Data Query / Usage (if applicable; optional fee)	If end-users or applications pay per query or heavy data usage, the fee could be very low per query	Distributed to maintainers of that data (e.g. indexers or the stakers of relevant data get a cut as "royalty" for usage). For instance, if someone frequently queries a certain claim, a portion of their query fees might flow to those who staked on that claim (encouraging staking on data that is frequently read, not just written).	A portion of fees is distributed to node operators and the treasury or burn to offset the network's maintenance costs.

Table 4: Illustrative Fee Distribution for Various Actions in Intuition.

In each case, notice how **everyone gets a piece of the pie** in proportion to their contributions:

- The actor always converts most of their fee into something of value for themselves (stake in data).
- Those who came before get a reward, which makes them whole for sharing the opportunity.
- The protocol (and by proxy all **TRUST** holders) accrues value to fund the system and token value.

These microtransactions at scale create a continuous incentive to contribute and curate. If Intuition achieves large usage (imagine millions of claims and identities being created and staked upon), even tiny fees can generate substantial revenue for the ecosystem. And because a share goes to token holders/treasury, the **TRUST** token is directly tied to platform activity – higher activity means more value flowing to **TRUST** (either through direct distribution or enabling the DAO to buyback tokens or invest in growth).

Incentive Compatibility and Game Theoretic Considerations

The token model is designed to be **incentive compatible**, meaning each participant's best strategy (to maximize their payoff) also tends to improve the network:

- If you want to earn a lot of rewards, you need to stake on information that will be widely accepted and used – thus you help highlight good information.
- If you want your token's price to rise, you should participate in governance to make decisions that increase adoption and utility – aligning individual profit with collective progress.
- If you want to cheat or submit false data, you face economic deterrents: you'll lose fees, and others will counter-stake against you (or in a governance dispute, you could be slashed), meaning the attempt is costly with low chance of profit.
- If you want to freeload (not contribute but use the data), you can to an extent (anyone can read the open data freely), but if you extract too much without contributing (e.g., building a closed-source app that monetizes Intuition's data without giving back), governance could introduce usage fees or limits that ensure major beneficiaries pay in. The open nature means this would be handled carefully, but it's a known consideration (similar to how open-source projects sometimes seek corporate sponsors for heavy use).

Market Dynamics: Because **TRUST** has various uses (staking, fees, governance), its demand

is bolstered by multiple sources:

- **Utility demand:** end users and developers needing **TRUST** to potentially pay certain fees or wanting it to stake.
- **Governance demand:** those who want influence will buy and lock **TRUST**.
- **Speculative/Value demand:** as Intuition's adoption increases, investors may purchase **TRUST** anticipating future demand.

By diversifying how **TRUST** is used, the token model avoids reliance on a single source of value. It's not *just* a governance token, or *just* a utility token; it's a hybrid designed to capture the upside of the entire platform's economy.

Standard Convergence Incentive: A subtle but critical incentive is the push towards data standardization. As discussed, if two competing schemas exist, the one that garners more usage will economically outperform the other, rewarding its backers. This tends to snowball: once a standard is winning, everyone jumps on it to earn rewards, and the losing standard loses whatever support it had. While it might sound harsh, this is how the network ensures *convergence* to one way of representing a concept, which is vital for a global knowledge graph. Importantly, those who correctly predict which standard will win (or help it win by advocating and staking) will earn more. This is a cooperative game dynamic where coordination is rewarded. The result is a more interoperable dataset and less fragmentation – an outcome that benefits all users of Intuition (the data is more unified) while economically benefiting those who coordinated.

Sustainability and Iteration

A well-aligned system should be **self-sustaining**. In Intuition:

- The continuous influx of new data and users paying small fees keeps rewarding existing contributors and funding the treasury.
- As emissions (inflation) decrease, the reliance shifts to **real usage revenue** as the main rewards source. Ideally, by the time token emissions slow down, the platform has enough usage (queries, enterprise integrations, etc.) that fees can meaningfully reward participants. This transition from “inflation subsidy” to “fee-driven revenue” is key for long-term sustainability.
- The DAO can adjust any part of this model if something isn't working. For example, if

some incentive is leading to unintended behavior (say people gaming IQ points or some niche of data is over-incentivized leading to spam), governance can dial it down. The economic design is not static; it's governed and upgradable.

Risk-Reward Balance: Participants face certain risks – a staker could lose potential rewards if they back a claim no one else finds useful, a token holder could see dilution if they don't participate while others bond and accumulate more **TRUST**, etc. However, these risks are transparent and manageable. They encourage active engagement (if you simply hold tokens and do nothing, your influence might dilute relative to those who lock tokens – but you can always choose to lock too). The system intentionally rewards activity and alignment (work, contribution, or commitment) over passivity.

To conclude this section, Intuition's economic incentives are crafted to ensure that **the rational, self-interested actions of individuals lead to a robust, growing knowledge network**. By tying data quality to financial rewards, and platform success to token success, Intuition aligns human and economic incentives with the production of trusted knowledge – a powerful synergy of blockchain and crowdsourced wisdom. In the next section, we'll discuss how the platform captures value and generates revenue, complementing these incentive structures.

Platform Revenue and Value Capture

A critical aspect of any token economy is how the platform generates **revenue** and how that value is captured by the token. While we have touched on fees in earlier sections, here we will systematically outline Intuition's revenue model and how **TRUST** holders benefit from it. We will cover the sources of revenue, the fee structure in detail, and the mechanisms by which revenue flows into the token economy (either via distributions, buybacks, or other means).

Revenue Generation Pathways

Intuition, as a decentralized data platform, has multiple potential revenue streams:

1. **Transaction Fees:** As previously described, actions like creating identities, making claims, staking, or otherwise interacting with the Intuition knowledge graph incur micro-fees. These can be considered analogous to “gas fees” for using Intuition's services.

Each of these fees is small (to encourage usage) but non-zero (to prevent spam and fund the system). Summed over large activity volumes, they constitute a primary revenue source.

2. **Data Access Fees:** While reading the raw data might be free (since data is on a public ledger), Intuition could offer value-added services for accessing data. This could include API endpoints for complex queries, analytics, or real-time updates. Enterprise users or heavy API users might pay subscription fees or per-call fees for these enhanced services (similar to how The Graph charges query fees to subgraph consumers). These fees can be denominated in **TRUST** or another currency but ultimately provide value to the ecosystem that can be captured by token holders (either by requiring payment in **TRUST** which creates buy pressure, or by collecting fees into the treasury).
3. **Premium Features and Services:** Intuition might introduce premium offerings such as verified badges for identities (which might require a fee and perhaps staking **TRUST**), or specialized curation services (like data quality audits, bespoke ontology support, etc.). These would primarily target organizations or power users who derive significant benefit from Intuition's data and are willing to pay for extra assurance or features. Revenue from such services would likely flow to the Intuition treasury or be shared with community curators involved in providing the service.
4. **Network Partnerships and Integrations:** Intuition could partner with other platforms (identity providers, DeFi protocols for credit scoring, AI companies for feeding verifiable data) in arrangements that generate revenue. For instance, an identity verification partner might pay Intuition a fee for every verification done using Intuition's data. Or a DeFi protocol might share fees if they use Intuition's reputational scores to enhance lending decisions. These are more speculative but represent how Intuition could monetize its role as foundational data infrastructure.
5. **Treasury Investments:** If the DAO treasury accumulates assets (from fees or bonding or initial allocation), it can deploy those assets productively. This might include providing liquidity in **TRUST** trading pairs (earning swap fees), investing in ecosystem projects (earning returns), or yield farming on stable reserves. The profits from these treasury activities effectively become revenue for the DAO, which can then redistribute or burn as it sees fit.

Out of these, the most direct and mechanically built-in are the **transaction fees** on the platform itself. They are predictable and directly tied to usage. Therefore, designing their structure properly is vital.

Token Value Support: Burn and Floor Mechanisms

Some token models implement a burn (destruction of tokens) to shrink supply as usage grows – famously, Binance Coin (BNB) or Ethereum’s EIP-1559 burning gas fees. Intuition could consider a partial burn model:

- For example, of the treasury’s share of fees, some portion could be automatically burned (say half of all ETH fees collected used to market-buy **TRUST** and then burn it). This creates a **deflationary pressure** proportional to usage. In a scenario of high usage, even if emissions are releasing new **TRUST**, the burn can offset some or all of that, achieving a net neutral or net deflationary supply. Such a mechanism provides a price support floor – basically the network “consumes” **TRUST** as it is used, similar to how increased usage of Ethereum now causes ETH to be burned, benefiting ETH holders.

Alternatively, the DAO might set a policy like “target a reserve ratio” – e.g., ensure the treasury holds assets equal to at least 25% of all **TRUST** market cap, through continuous bonding and buybacks, which gives a kind of book value backing to each token. This was the concept in protocols like Olympus DAO (where each token was backed by treasury assets). While Intuition’s focus is more on utility than creating a reserve currency, having a strong treasury provides confidence and can mitigate downside (the DAO can intervene if token price crashes far below the value of assets per token, by buying back aggressively).

One could imagine down the road Intuition being **self-funding**: the revenue from platform usage could pay for ongoing development (via grants to devs), community events, and so on, without needing external capital – all authorized by **TRUST** governance. This is the ideal end-state of a decentralized project: it becomes **a self-sustaining economy**, where token holders govern a treasury that can continuously improve and expand the ecosystem, creating a positive feedback loop between platform success and token value.

Conclusion

In conclusion, TRUST operationalizes a normative principle that the informational substrate of the internet—data, digital identities, and the mechanisms of discovery—ought not reside under the exclusive dominion of any centralized entity. Instead, ownership and control are distributed across all participants through an explicit, token-weighted graph. Each assertion in

Intuition's graph is collateralized by staked TRUST, rendering veracity a quantifiable and economically contestable property, while simultaneously precluding unilateral appropriation of the corpus.

The resulting incentive architecture is double-layered. On the epistemic plane, actors are rewarded for the timely provision and verification of useful data; on the economic plane, they internalize the negative externalities of misinformation via slashing and stake decay. By pricing credibility rather than attention, the system aligns individual profit maximization with the collective objective of data integrity. Empirically, this framework transforms the knowledge graph into a self-refining public good whose marginal improvement is financed endogenously by its users.

Governance proceeds through token-weighted deliberation, enabling continuous protocol evolution and the calibration of bonding curves, fee redistributions, and operator-level security parameters without reinstating centralized veto power. Because stake accrues heterogeneously to identities and claims, no single entity can capture disproportionate rents; rather, value distribution asymptotically approaches the contribution profile of the network's participants. This establishes a credible commitment to pluralistic ownership while maintaining the adaptability required for long-term system viability.

In summary, TRUST advances a cryptoeconomic model in which the stewardship of digital knowledge is both shared and incentive-compatible. It converts truth into an investable, excludable, yet non-rivalrous asset and endows the wider community with the formal governance apparatus necessary to sustain it. The paradigm is thus not a rhetorical aspiration but a mathematically specified pathway toward a decentralized epistemic infrastructure whose evolution will be authored collectively—one stake at a time.