

PROOF OF CONTRIBUTION



TALVORY

The Knowledge Layer for Machine Intelligence

Human contribution, verified and licensed as
data assets. Settlement returns to the people
who produced the corpus.

TLVY
Technical and Economic Framework
Version 1.0 · Pre-Issuance

TABLE OF CONTENTS

1.0 Executive Summary — The Missing Layer

2.0 Market Diagnostic: The Intelligence Asymmetry

3.0 Protocol Architecture

4.0 Layer I — The Commons: Proof of Contribution

5.0 Layer II — The Agent Economy

6.0 Layer III — The Sovereign Compute Mesh

7.0 The TLVY Token: Multi-Layer Utility and Value Capture

8.0 Supply Architecture: Distribution, Release, and Deflation

9.0 Issuance Framework

10.0 Governance, Compliance, and Data Sovereignty

11.0 Roadmap, Organization, and Disclosures

Appendix A — Definitions and Terminology Index



Abstract

Frontier model capability is no longer constrained primarily by compute. It is constrained by the availability of licensed, provenance-bearing human knowledge in languages, domains and economic contexts that have never been systematically collected. Existing collection mechanisms are either unconsented, unverified, or economically extractive toward the people who produce the underlying material.

Talvory is a coordination protocol for producing, verifying, licensing and settling that knowledge. Contributions are classified, peer-reviewed and bonded; verified output is registered as a licensable data asset with a defined consent scope and a revenue waterfall back to its producers. Autonomous agents trained on the resulting corpora operate as the demand surface, generating the usage that funds the supply side.

TLVY is the protocol's bonding, metering and governance asset. It is required for validator security, licence settlement between counterparties with no prior contractual relationship, and the parameter governance of a language registry that cannot be responsibly administered by a single foundation.

This document states what the protocol does, how its economics close, and – in Chapter 7 – which of its functions would continue to work without a token at all.

1.0 Executive Summary – The Missing Layer

1.1 The Thesis in One Page

The economics of frontier AI have inverted. Through 2023 and 2024 the binding constraint was compute; capital flowed into accelerators, and the organisations that controlled large clusters controlled the frontier. That constraint has loosened faster than most forecasts anticipated. The next constraint is upstream and less tractable: the stock of high-quality, human-generated, provenance-clear material available for training is finite, and a growing share of new material is itself machine-generated.

Several large model developers have publicly acknowledged approaching or exhausting readily available high-quality web text, and multiple studies have documented measurable degradation when models are recursively trained on synthetic output.

Synthetic data extends the runway. It does not extend it indefinitely, because synthetic generation compresses the diversity of the distribution it imitates; each recursion loses tail coverage first. The long tail is precisely what the current generation of models handles worst and what the next generation of products – clinical decision support, public-sector deployment, regulated financial services – requires most.



That tail is not scarce because it is hidden. It is scarce because there has never been a mechanism to pay the people who hold it. Roughly seven thousand living languages are spoken today; fewer than a hundred have meaningful coverage in publicly documented training corpora. The same asymmetry holds in professional domains within high-resource languages, where the material that matters is proprietary, private, or held in formats no crawler can reach.

Talvory is infrastructure for the collection, verification and licensing of that material, coupled to a demand-side product that generates usage. Its commercial claim is that licensed, provenance-bearing, jurisdictionally clean corpora will carry a durable price premium over scraped equivalents — first in regulated procurement, eventually everywhere.

1.2 What Talvory Is, and What It Refuses to Be

The protocol has three layers and one asset.

Layer I, the **Commons**, converts human contribution into verified data assets. Contributors submit work in five classes: linguistic material, domain annotation, evaluation judgement, translation pairs, and structured local knowledge. Each submission enters a review process with a challenge window; validated output is deduplicated, registered, and assigned a consent scope governing onward use. Licences are issued against the registry; settlement flows back through a defined waterfall.

Layer II, the **Agent Economy**, is the demand surface. Autonomous agents execute defined work on behalf of micro and informal enterprises: order intake, reconciliation, inventory signalling, receivable follow-up, catalogue and content generation. These agents run on models fine-tuned against the Commons, which means agent accuracy is a function of Commons depth. Every executed task returns structured evidence of what worked and what failed, and that evidence is itself a contribution.

Layer III, the **Compute Mesh**, aggregates inference capacity under residency constraints. It exists because a material share of the demand is legally unable to route data across certain boundaries, and because the cost differential between aggregated regional capacity and centrally priced inference is the margin that funds the Layer II price point.

TLVY is the asset that binds these layers.

It is easier to state what the protocol declines to be, because those decisions are the ones a reviewer will test first.

Talvory does not compete on raw compute. The mesh is a cost instrument and a compliance instrument, not a claim of superiority over hyperscale infrastructure. Aggregated regional capacity wins on residency and on price at moderate utilisation; it does not win on peak throughput for frontier-scale training, and the protocol does not attempt to.



It does not operate a general-purpose crawler. Every unit in the Commons carries an identifiable contributor, a recorded consent scope and a verification history. This makes the registry slower to build and materially easier to license. The trade is deliberate.

It does not sell an agent product as its primary business. Agent subscriptions are the mechanism by which usage, and therefore data demand, is generated at scale. A reviewer who evaluates Talvory as a vertical SaaS company will find the unit economics unremarkable; that is the correct reading of the subscriber layer in isolation and the wrong reading of the system.

Network compensation is purely performance-denominated: emissions and fee distributions compensate verifiable data provisioning, validation throughput, and compute availability. The protocol does not offer passive yield or fixed financial returns; all economic compensation requires active cryptographic or computational proof of work.

1.3 The Value Capture Statement


A token that sits alongside a protocol without touching it is a liability, and the market has become efficient at identifying them. TLVY is attached at three points where the alternative is either structurally unavailable or economically inferior.

Bonding. Validators, compute nodes and agent operators post TLVY as a performance bond. Verification without confiscatable stake degrades to reputation, and reputation has no cost of falsification.

Settlement. Licence fees and cross-entity work settlements occur between parties with no prior contractual relationship, frequently across payment-rail discontinuities. TLVY denominates those settlements and is converted at the gateway boundary; the burn allocation on conversion is a function of protocol utilisation rather than a discretionary decision.

Governance. A language registry that admits new corpora, sets quality floors and determines revocation policy cannot be administered by a single foundation without becoming a political object. TLVY weights that constituency. Chapter 10 specifies which decisions are permanently outside its reach.

The mechanism connecting these is described in 7.2 and 7.3. The supply consequences are modelled in Chapter 8. Before either can be assessed, the market condition that makes the protocol necessary has to be established — which is the subject of the next chapter.



2.0 Market Diagnostic: The Intelligence Asymmetry

2.1 The Data Wall and the End of the English Corpus Advantage

The supply of high-quality training text behaves like a depleting stock rather than a growing one, for three reasons that compound.

The first is simple exhaustion. Indexed public web text is finite, and the highest-quality portion of it — edited prose, technical documentation, refereed literature, well-moderated discussion — was crawled early. What remains is disproportionately low-signal, templated, or commercially gated.


The second is contamination. As generative systems produce an increasing share of newly published text, the ratio of human-authored to machine-authored material in the crawlable pool deteriorates. Recursive training on such material degrades distributional coverage, and the degradation appears in the tail first. Benchmarks built on common distributions can mask this for several training cycles.

The third is legal. The permissive era of unlicensed scraping is closing unevenly



across jurisdictions. Procurement rules in regulated sectors increasingly require documented provenance for training material, which converts an informal cost advantage into a formal cost for anyone who cannot produce that documentation. Provenance requirements will spread from public-sector procurement into enterprise contracting within the planning horizon of this document. The falsification test is straightforward: if large enterprise buyers continue to accept undocumented training data without price or compliance penalty, the premium the protocol depends on narrows.

Talvory is built for the third constraint, which is the durable one. Compute constraints are solved by capital. Legal constraints are solved by testimony.



2.2 The Representation Gap and the Value Extraction Paradox

Two quantities are routinely conflated in discussions of language AI: how many people speak a language, and how much value is currently served in it.


The distribution of training material across languages is far more skewed than the distribution of speakers. A small set of languages accounts for the overwhelming majority of indexed text, while the majority of living languages have effectively no representation.

The second-order effect is the one with commercial consequences. Where representation is absent, capability is absent, and where capability is absent, the market is served either by nothing or by degraded outputs that users compensate for manually. That compensation is invisible in any published statistic, which is why the market appears smaller than it is.



The extraction side of the same phenomenon runs as follows. Material is produced locally, at local cost, in a form that has local value. It is collected without payment, aggregated elsewhere, and returned as a product priced for high-income markets. The producer captures nothing and pays for the output. This is not a moral observation that requires a remedy on moral grounds; it is a pricing observation, and it predicts that the first party to offer payment for the material will face materially lower collection costs than the last party to take it for free.

Contributors respond to payment incentives in the same direction as other economic actors, and will supply verifiable work when the terms are legible and the compensation settles reliably. The falsification test is contributor retention after incentive taper, which the phase gates in Chapter 11 require as a measured quantity.



2.3 The Missing Layer in Sovereign AI Stacks

A number of states have begun funding AI capacity as public infrastructure: compute programmes, national model efforts, talent pipelines, and in several cases sovereign investment vehicles with mandates to underwrite both. The pattern is consistent enough to treat as a structural feature rather than a series of coincidences.

What these programmes systematically under-provision is the data layer.

Compute is procurable. A procurement office can specify accelerators, benchmark them, and award a contract. Models are procurable in the same fashion, or licensable. The material required to make either useful in a given language and administrative context is not procurable in that form, because at the point of purchase it does not exist as an inventory. It exists as latent capacity distributed across a population.



The consequence is that sovereign programmes routinely fund the two layers that can be bought and leave the layer that must be produced. The result is capacity without local adequacy: infrastructure that runs competently in a handful of high-resource languages and degrades in the languages its own constituents use. Remedying this through internal hiring is possible at small scale and prohibitive at the scale the condition requires – the volumes involved are in the range of hundreds of thousands of contributor-hours, spread across hundreds of languages and dialects, in workflows too varied for a centralised workforce.

This is the structural gap the protocol addresses, and it is worth stating its limits precisely. A protocol does not replace a national programme and does not compete with one. It supplies the input those programmes require and cannot internally generate. The commercial consequence is a counterparty that is not another data network, and a procurement motion that does not run through developer adoption at all.

At least some sovereign and quasi-sovereign AI programmes will contract for external corpus supply rather than building collection capacity internally. The falsification test is observable within the first two phases: if no programme contracts, the Layer I revenue model rests entirely on private-sector licensors, which reduces the modelled licence revenue by the share attributed to public counterparties in Figure 4.4.

2.4 The Anchor Market Profile

Deployment begins in a single market, selected against six criteria rather than chosen. The criteria are stated in full because the selection is the load-bearing strategic claim of the document, and because a reader who disagrees with the choice can substitute their own weighting and see whether the conclusion survives.


Criterion	Why it matters	What disqualifies a market
Linguistic fragmentation	Determines the size of the unserved corpus and the collection surface	Monolingual markets offer no marginal collection value
Micro-enterprise density	Determines the depth of Layer II demand and the volume of task-generated evidence	Economies where formal enterprise dominates offer few agent-deployable workflows
Mobile payment rail maturity	Determines whether Layer II can be billed without building a collection stack	Cash-dominant markets raise the cost of small-ticket billing above its value
Policy momentum in AI capability	Determines whether Layer I licence demand includes a public counterparty	Absence of any state AI programme removes the highest-margin licence class



Criterion	Why it matters	What disqualifies a market
Mobile and network penetration	Determines whether contributors can participate without hardware subsidy	Populations without device access cannot enter the Commons
Cost of agent delivery	Determines whether Layer II margins can fund Layer I collection	High inference, connectivity and support costs invert the model

A market that scores highly on fragmentation, micro-enterprise density and payment-rail maturity while retaining low delivery costs will exhibit a materially better collection-cost-to-licence-value ratio than a market that scores highly on policy momentum alone. This is the assumption the Anchor selection rests on. It is testable in Phase I and its failure is recoverable, because the collection mechanism is portable across markets in the same archetype while the licence relationships are not.

The Anchor is referred to in this document as a high-density maritime emerging market archetype. Using the archetype rather than the identity serves a purpose beyond discretion: it forces every downstream claim to be stated as a property of the archetype, which means those claims can be evaluated against other markets that share the properties. A reader evaluating a different geography can test the model directly against the six criteria without needing to discount for local detail they cannot verify.





3.0 Protocol Architecture

The protocol exists to answer one question: how does unverified human work become a licensable asset with a provenance record strong enough to survive procurement scrutiny, without routing through an intermediary whose interests diverge from both sides of the transaction?

Every structural choice in this chapter follows from that question. Where a design decision appears arbitrary, it is usually the residue of a constraint described earlier in this document or one that the implementation phase will surface.

3.1 Five Design Principles

Verification precedes value. Nothing is paid on submission. A contribution enters provisional credit, and provisional credit is an accounting entry, not a transferable balance. Settlement occurs after adjudication. The delay is the mechanism: it removes the incentive to flood the queue with low-cost submissions whose marginal cost to the submitter is near zero.

Bonds secure what reputation cannot. Reputation scoring is useful for ranking and useless for enforcement. A validator with an accumulated score and no forfeitable position faces no cost when the score is wrong, because the worst outcome is a deduction against a quantity that cannot be spent. Adjudicators therefore hold bonded positions, and the bond is denominated in TLVY. Reputation determines how much bond a validator must post to reach a given volume of adjudication; it does not substitute for the bond.

Consent is recorded as a scope, not as a copy. The chain carries a claim about what may be done with a corpus — which licence classes, which jurisdictions, which retention limits, and until which date. It does not carry the corpus. This distinction has a practical cost: every licence issuance requires a scope-match check against the registry, and a scope mismatch blocks the transaction rather than generating a warning. The alternative, storing material on-chain for the sake of auditability, produces a permanent record that no contributor can withdraw and no regulator can accept.

Demand and supply are one loop rather than two businesses. Layer II agents exist because a data network with no usage has no way to discover which corpora matter. Every executed task tells the Commons something about where model output failed in a real workflow, and that failure signal is more valuable per unit than bulk collection because it is conditioned on a task someone was willing to pay for. An agent product evaluated on its own subscriber economics looks unremarkable; the correct unit of analysis is the joint system.

Finally, **permission boundaries** are made visible rather than smoothed over. Some protocol components are permissionless, some require bonded participation, and some are legally gated by the jurisdictions in which they operate. A design that presents all three as uniform will attract counterparties who discover the



difference during diligence, which is a worse outcome than disclosing it in advance.

The fourth principle is the one most likely to be challenged by investors who prefer separable business units. It is retained because separating them would eliminate the feedback path that makes Layer I collection targeted rather than speculative.

3.2 The Three-Layer Stack and Permission Boundaries

Each layer carries a distinct function and a distinct metering surface.

Layer I meters contribution. The unit is a verified contribution record, priced by class and quality tier. Layer II meters execution: the unit is a completed task, priced per task or bundled into a subscription. Layer III meters capacity, priced per unit of inference or per reserved throughput window.

The permission classes cut across all three.

Component	Permission class	Rationale
Contribution submission	Permissionless	Open entry is what makes the collection surface wide enough to matter
Peer review pool	Reputation-gated	Reviewers accumulate accuracy records; low-accuracy reviewers lose routing priority
Adjudicator validation	Bonded	Adjudication carries economic consequences and must be insured
Compute node operation	Bonded, with residency attestation	Residency claims require something forfeitable behind them
Agent deployment	Permissionless, with operator bond above a volume threshold	Low-volume experimentation should be free; sustained commercial operation should not
Licence issuance	Contractual, gateway-mediated	Licence terms are legal instruments and cannot be reduced to an on-chain function call
Settlement	Legally gated	Executed through regulated intermediaries; the protocol issues instructions and does not hold funds

The classification follows from existing financial and data-protection regimes in the jurisdictions where the protocol intends to operate, which distinguish between holding, transmitting and instructing.

The practical consequence for TLVY is that its utility concentrates at the bonded



and metered boundaries. Contributors do not need TLVY to submit work; they receive it. Agent subscribers do not need TLVY to use an agent; they pay in local currency through a gateway, and the protocol's fee accrues in TLVY at the conversion step. A design that required end users to acquire and hold a token before accessing a service would fail both the usability test and, in a number of markets, the regulatory one.

3.3 Actor Topology: Six Classes, Three Trust Edges

Six participant classes interact with the protocol, and the trust relationship between any two of them falls into one of three categories.

Contributors submit work and hold no bonded position. They bear no downside beyond the rejection of a submission, which is appropriate given that their capital at risk is their time.

Adjudicator validators adjudicate and carry forfeitable bonds. They also hold governance weight in the chamber that sets quality floors, which creates the obvious conflict: a validator class with seats in the body that determines admission standards can raise barriers to entry against competitors. Chapter 10 addresses this through seat-term limits and an independence requirement on the technical chamber, but the tension is structural and will not be engineered away entirely.

Compute nodes provide inference and training-shard capacity under residency attestation. Their bond is sized against the throughput they have declared, not against their reputation, because an overclaimed throughput declaration is the failure mode that matters.

Agent operators deploy products against the Commons and pay per-task fees. They are the only class whose incentive runs directly against collection cost — a cheaper Commons is better for them and worse for contributors — which is why their fee flows into the contributor pool rather than being retained.

Data licensees hold no on-chain position at all. They are contractual counterparties, and their trust relationship with the protocol is legal rather than cryptographic. This is deliberate: forcing enterprise licensees into a token relationship would add friction to the highest-margin revenue line for no security benefit.


Settlement gateways are regulated intermediaries operating under their own licences. Their relationship with the protocol is contractual, and the protocol's exposure to their conduct is a counterparty risk documented in the risk register.

The third edge type — trust-minimised, bonded, contractual — maps onto these classes without overlap. On-chain verification handles contribution provenance, adjudication records and supply accounting. Bonded edges cover everything where a party could profit from misreporting their own conduct. Contractual edges cover everything the protocol cannot see.



3.4 The Compounding Flywheel

Task execution produces a failure record. That record — which request the model misparsed, which field it filled incorrectly, which local term it did not recognise — enters the Commons as contribution-class material of the evaluation type. Adjudicator verify it, the registry admits it, a fine-tuning run consumes it, and agent accuracy improves on the specific workflow that generated the failure.




The loop closes in the direction that matters: usage generates the training signal that improves usage.

Two conditions can break it, and both are measurable rather than theoretical.

The first is verification latency. If the interval between submission and settlement exceeds the contributor's tolerance, submission volume falls regardless of the nominal reward, because contributors with genuine local expertise have alternate uses for their time. Phase I therefore treats median time-to-settlement as a gate criterion rather than a performance metric.

The second is the depth threshold. Agent accuracy improves as a function of Commons depth only above a coverage level where the fine-tuning corpus carries enough density in the target domain to move behaviour. Below that level, additional collection produces measurably nothing. The threshold is reachable within the Phase II envelope for the primary language and two or three adjacent ones, and not for the full language set, which is why the capability ladder in 5.3 sequences by domain rather than by language count.



4.0 Layer I – The Commons: Proof of Contribution

4.1 Contribution Taxonomy and the Proof of Contribution Consensus

Contributions arrive in five classes, and the class determines how the work is reviewed. Treating them uniformly would either over-constrain the verifiable classes or under-constrain the subjective ones.

Linguistic material covers speech recordings, transcribed text and dialect variants, with a required metadata envelope describing speaker provenance, recording conditions and regional register. Its verification method is largely mechanical: signal quality checks, transcript agreement across two independent transcribers, and duplicate detection against the existing corpus.

Domain annotation covers labelled entities, classification and structured extraction over domain-specific material. Verification is by sampled re-annotation and inter-annotator agreement against a quality threshold that varies by domain difficulty.

Evaluation judgement covers pairwise ranking, preference elicitation and error



identification over existing model outputs. This is the class whose value is least legible and whose verification is hardest, because the ground truth is a human judgement rather than an observable property. It is reviewed through agreement statistics across a panel rather than by adjudication against a standard, and it is also the class that generates the flywheel signal described in 3.4, which makes it disproportionately important relative to its volume. Formally, consensus convergence over subjective annotations is evaluated via an empirical multi-annotator agreement coefficient:

$$\kappa = \frac{P_o - P_e}{1 - P_e} \geq \kappa$$

$$P_o = \frac{1}{|C|} \sum_{i \in C} \mathbb{I}(a_{i1} = a_{i2})$$

$$P_e = \sum_{k \in K} p_{k1} \cdot p_{k2}$$

Parameter Specifications:

- P_o denotes the empirical observed agreement ratio across independent peer reviewers for corpus unit C .
- P_e denotes the baseline expected probability of stochastic chance agreement across domain category space K .
- $\kappa_{\text{threshold}} \in [0.65, 0.85]$ serves as the parametric gate: contributions enter Validated status exclusively upon satisfying this convergence boundary within the active challenge window.

Translation pairs and parallel text are verifiable through round-trip checking and through consistency scoring against a reference subset.

Structured local knowledge — procedural, regulatory, agricultural, customary and administrative material rendered as question-and-answer or decision-tree form — is verified by subject-matter review, which is the most expensive verification method available and the reason this class carries the highest tariff.

Class assignment happens at submission and is contestable. A contributor who believes their submission was misclassified can request reclassification during the challenge window, and a successful reclassification carries the difference in tariff. Misclassification is a predictable failure mode when contributors are learning the taxonomy, and treating it as fraud rather than error would suppress exactly the submissions the network most needs.

The consensus path runs through four states. A submission enters as **Submitted**, is assigned to the peer review pool, and on receiving the required number of




independent reviews passes to **Provisionally Credited**, at which point the tariff is escrowed but not transferable. An Adjudicator validator then either confirms the provisional outcome or escalates it. Confirmation moves the record to **Validated**, and the challenge window opens. Settlement — the point at which credit becomes transferable and the contributor can withdraw — occurs when the window closes without a sustained challenge.

Rejections are classified as repairable or non-repairable. Insufficient metadata, contaminated audio and missing consent fields are repairable, and the submission is returned with the specific deficiency identified. Fabricated provenance, duplicate submission under altered metadata, and consent scope violations are non-repairable and carry a reputation penalty. The distinction exists because most rejected submissions are defective rather than fraudulent, and a system that treats both identically trains contributors to stop submitting after their first mistake.

Proof of Contribution

End-to-end consensus flow · Five swimlanes · Asset lifecycle state band



Escrow behaviour shown explicitly. Failure branches separate repairable from non-repairable defects. Every state carries a defined exit condition. Challenge window is fixed at 48h; during this period a bonded challenge can trigger re-adjudication. Settled state is the only terminal success state.

4.2 The Adjudicator Verification Network

Validator selection begins with a competence attestation in a specific language or domain pairing. Attestation is earned through scored adjudication against a held-out set, and it degrades if the validator's adjudications diverge from panel consensus over a rolling window. A validator competent in one low-resource language carries no admission to another.

Bond requirements scale with adjudication volume and inversely with recent accuracy. A new validator posts a baseline bond for small volumes; sustained accuracy unlocks higher throughput at a lower bond-to-volume ratio. This produces a system where the most reliable validators are also the cheapest to operate, which matters because validation cost is the largest variable cost in Layer I.



Slashing has three destinations rather than one. False validation returns the bond to the harmed party — the contributor whose work was wrongly rejected, or the licence counterparty who received a misrepresented corpus. Duplicate or fabricated adjudication is burned, since there is no identifiable harmed party. Procedural breach without misrepresentation routes to the treasury. The three-way split exists because routing all penalties to a single destination produces the wrong incentives at the margin: validators facing a shared penalty pool treat slashing as a cost of doing business rather than a transfer to a specific victim.

The network has an unsolved constraint that should be stated plainly. Low-resource languages are low-resource partly because qualified validators are scarce, and a protocol that requires expert adjudication before settlement cannot grow faster in a language than its validator pool allows. Three mitigations are in scope: tiered validation, where mechanical checks resolve objective classes without Adjudicator involvement; cross-language panel review for closely related varieties, which extends coverage without claiming validator competence that does not exist; and a validator development programme funded from the treasury, which converts contributor reputation into validator capacity over time. None of these removes the constraint. They change its slope.

Mechanical verification is sufficient for the linguistic material class at the volumes the first two phases require, and Adjudicator involvement can be reserved for annotation, judgement, translation and structured knowledge. If the mechanical pass rate for linguistic material falls below the modelled figure, Layer I cost rises against Figure 4.4's waterfall and the answer is a narrower class scope rather than a higher tariff.

4.3 Data Assetization and the Licensing Engine

Verified contributions aggregate into corpora through normalization, deduplication and consent-scope reconciliation. Scope reconciliation is where most aggregation attempts fail: combining two corpora with incompatible consent scopes produces a derived asset that no licence can lawfully cover, because the intersection of the scopes is narrower than either original. The registry computes that intersection explicitly and refuses to register an asset whose intersected scope is empty.

Licences fall into three tariff classes. Research and non-commercial use carries the lowest tariff and the narrowest redistribution rights. Commercial model training permits derived weights but prohibits redistribution of the corpus itself. Commercial redistribution permits onward sale and carries the highest tariff plus a reporting obligation on downstream use.

Sensitive categories — material touching health, financial behaviour, or identifiable personal attributes — are licensed through compute-to-data rather than transfer. The licensee submits a training or evaluation job to a controlled environment inside the residency perimeter, receives aggregate results, and never obtains the corpus. This is slower and more expensive for the licensee, and it is the only



arrangement under which several categories can be licensed at all.

Revocation requires an honest statement of what is achievable. A consent withdrawal can remove a corpus from future licensing, trigger deletion from active storage, and prevent any new derived asset from being registered against it. It cannot recall model weights trained on the corpus before withdrawal, and no cryptographic mechanism makes that possible. Consent scopes therefore specify retention and training-use terms at the point of contribution, so that contributors consent to a defined downstream use rather than to an unbounded one, and revocation governs the future rather than pretending to govern the past. Compensation on revocation is available from the treasury where the original licence has already settled.


Enterprise procurement interfaces integrate with standard off-chain data consortia through authenticated API endpoints. Downstream model developers (e.g., enterprise SLM fine-tuning pipelines) settle bulk licensing in fiat via compliant clearing corridors; programmatic conversion and protocol burn occur deterministically via the Settlement Gateway Layer at execution time, anchoring token deflation directly to verified external enterprise consumption rather than speculative float.

4.4 The Revenue Waterfall

Gross licence fees cascade through six stages before any residual reaches the protocol's discretionary balance.

Gateway and settlement costs are deducted first, because they are payable to third parties regardless of the fee's size. The protocol treasury allocation follows, funding operations, audits, the validator development programme and legal defence. The validator pool is paid next, sized against adjudication volume rather than licence value, which keeps verification cost attached to the work performed. The contributor pool receives the largest single allocation, distributed by contribution record class and quality tier. A burn allocation is applied on the TLVY conversion step rather than on the fiat fee, which is why the deflationary behaviour in 8.3 is a function of utilisation. Any residual sits in the treasury and is subject to the governance path in Chapter 10.

Splits shift by tariff class. Research licences carry a thinner treasury allocation and a wider validator allocation, since research corpora require more adjudication per unit of fee. Commercial redistribution licences carry the reverse. A licensee negotiating terms therefore negotiates against a published band rather than a bespoke arrangement, which shortens contracting and removes the most common source of counterparty dispute.



5.0 Layer II – The Agent Economy

5.1 The Micro & Informal Enterprise as the Unit of Demand

The demand surface for agent execution is defined by an operational threshold rather than by sector, geography or revenue band. A business sits below that threshold when the cost of acquiring, configuring and maintaining conventional enterprise software exceeds the value of the workflow it would automate. This is not a statement about capability or sophistication. Many such businesses run complex operations with real inventory, real receivables and real counterparties. What they lack is the margin per transaction to absorb implementation cost, seat licensing and training overhead.

Three structural features define the category.

Workflows repeat heavily but at low value per instance. An order confirmation, an inventory check or a payment reminder is worth a small fraction of a currency unit in labour saved. Software priced per seat cannot earn its cost against that transaction value, and software priced per instance has no mechanism to bill at that granularity without a payment rail.

Operating attention is discontinuous. The owner of such a business performs most functions personally and cannot dedicate sustained attention to any single system.



Interfaces that require a session, a login sequence or a configured dashboard lose to an interface that requires a sentence spoken between two other activities.

Devices are mobile-primary and often mobile-only. Any assumption of desktop access, stable bandwidth or peripheral hardware excludes a majority of the category.

Mobile payment rails have penetrated this category substantially faster than enterprise software has, in the Anchor archetype and in comparable markets. The consequence is that billing friction at the point of sale has largely been solved while the operational workflows upstream of the sale remain unsupported.

This is the gap Layer II occupies, and the framing determines what gets built. An agent delivered to this category cannot be a scaled-down enterprise product, because the constraint is not feature depth but the total cost of adoption measured in the owner's time. The relevant question is whether a task can be completed with fewer owner-seconds than the manual alternative, including the seconds spent learning the system.

A voice-primary interface reduces the learning cost of first use below the threshold at which trial converts to retention, and that this effect is larger than any achievable gain from improving task accuracy. The falsification test is trial-to-retention ratio measured separately for voice-primary and text-primary cohorts within the same operator.

5.2 Voice-First Intent and the Execution Pipeline

Voice is treated as the primary input surface rather than an accessibility feature. The reasoning is economic. A text interface requires the user to translate intent into a structured query, which imposes a cost on every interaction, and that cost is paid in the scarcest resource the user has.

Two consequences follow for the pipeline design.

Parsing must tolerate ambiguity rather than resolve it optimistically. Spoken intent arrives underspecified: quantities omitted, products referred to by local trade names, counterparties unnamed. A pipeline that guesses produces wrong actions, and a wrong order or an incorrect payable is more costly to the user than a delayed one.


Every stage that can commit the user to an obligation must have a confirmation gate above a defined confidence threshold. The threshold varies by consequence class, which is the design detail that separates this pipeline from a conversational wrapper. Generating a draft costs nothing if wrong. Sending a message to a counterparty costs reputation. Committing an order costs money. Adjusting a price costs margin.

The pipeline therefore runs in seven stages, and the seventh exists because the sixth is not sufficient.



Intent capture records the request in the user's language, preserving the original utterance alongside any normalization, since the raw form is itself contribution-class material. Intent parsing extracts an action type and a slot set, and flags every slot it could not fill. Task planning drafts a sequence against a capability template. Confidence evaluation scores the draft against the slot completeness and the historical accuracy of that template for that user. The confirmation gate presents only the ambiguous or consequential fields, on the principle that asking the user to verify a draft they already dictated trains them to stop reading confirmations. Execution calls the external systems the capability template is bound to. Reconciliation compares the external system's returned state against the planned intent, and where they diverge, the divergence is logged as a failure record rather than silently corrected.

The reconciliation stage is the one most often omitted in comparable systems and the one that makes the data flywheel function. An execution that returns an unexpected state is a labelled example of the model failing in a live workflow, produced at no marginal cost to anyone.



5.3 The Capability Ladder

Agent capability is sequenced in five levels, and the sequencing runs by workflow domain rather than by language coverage or model size.

Level one covers single-task assistive output: draft content, translate a catalogue entry, summarize an incoming message thread. The output is advisory, nothing is committed, and the failure cost is bounded by the user's willingness to reject a draft.

Level two covers record-keeping on behalf of the user: capturing transactions, maintaining inventory state, recording receivables. The failure cost rises because the record becomes the basis for later decisions, though no external party is affected.



Level three covers outbound communication to counterparties: follow-ups, confirmations, dispute responses. Reputational cost enters here, and this level is the first at which a confirmation gate is mandatory regardless of confidence score.

Level four covers transaction commitment: placing orders with suppliers, confirming delivery schedules, accepting settlement terms. Financial exposure begins, and this level is gated on operator bonding.

Level five covers cross-system fulfilment, where an agent coordinates across inventory, order, payment and logistics systems that were not designed to interoperate. This is where the category's operational value concentrates, and it is also where the failure modes are least diagnosable, which is why it sits last.

Advancing a domain from one level to the next requires measured accuracy at the current level over a defined task volume. The gate is per-domain, which means an operator serving the food service vertical at level four may simultaneously be serving the textiles vertical at level two. This is deliberate: an aggregate accuracy figure across all domains would permit a strong domain to carry a weak one into a failure class it has not earned.

Level four accuracy is achievable within the Phase II task-volume envelope for two verticals in the Anchor's primary language. Level five is presented as a Future Option.

5.4 Unit Economics

The Layer II model turns on a single variable: the share of tasks resolved without human intervention in the execution path. Every interaction that requires an operator to intervene adds cost measured in support minutes, and support minutes are priced in a labour market where the margin per task is small.

Four cost components sit below revenue per enterprise. Inference cost scales with task volume and is bounded by the mesh pricing described in 6.3. Orchestration cost covers the routing, template maintenance and tooling that the agent requires regardless of task count. Settlement gateway fees are charged per transaction and therefore fall as ticket size rises, which is a mild structural advantage for higher-value verticals. Support cost is the residual, and it is the component that determines whether the model works.


Component	Behaviour	Lever
Inference	Linear in task volume, bounded by mesh unit price	Node mix and batching
Orchestration	Semi-fixed, rises in steps with template count	Template reuse across enterprises in a vertical
Settlement	Per-transaction, declines as share of value	Vertical selection toward higher ticket sizes
Support	Nonlinear, concentrated in first 60 days per enterprise	Onboarding design; capability ladder gating



Support cost is nonlinear in a specific way: it clusters in the early period after an enterprise onboards, then falls to a low baseline until the enterprise's task mix changes, at which point it steps up again. A model that averages support cost across the customer lifetime will overstate steady-state margin and understate churn driven by a bad first month.

Distribution compounds the problem. Reaching enterprises in this category one at a time through paid acquisition does not clear the average revenue per enterprise. The channel structure that works is aggregative: distributors who already serve groupings of such businesses, trade associations, payment service providers with merchant bases, and local implementation partners who carry the onboarding burden in exchange for a revenue share. The share is expensive and it is cheaper than the alternative.

Subscription pricing is set against automation depth rather than seat count, and the ladder gates determine what depth an enterprise can be sold. Selling level four capability to an enterprise that has not cleared level two is the primary driver of early churn in agent deployments and is prohibited at the operator level rather than merely discouraged.





6.0 Layer III – The Sovereign Compute Mesh

6.1 Residency as a Commercial Requirement

The mesh exists because a material and growing share of the demand surface cannot lawfully route certain categories of data across certain boundaries. This is a procurement condition rather than a philosophical position, and it converts into revenue straightforwardly: a counterparty that is legally unable to use cross-border inference will pay for compliant inference at a premium, and the premium is bounded by the cost of the alternative, which for many such counterparties is not building anything at all.

Three workload classes are affected to differing degrees.

Public-sector and regulated-sector inference typically requires that input data be processed within a jurisdiction, with an audit trail establishing that it was. Where such requirements exist, they are usually absolute rather than risk-weighted, which means a non-compliant provider is excluded rather than penalised.

Training and fine-tuning on Commons corpora operate under the consent scopes described in 4.3, which may specify processing boundaries independently of where the contributor resides. A corpus contributed under a scope restricting processing to a defined perimeter cannot be consumed outside it, and the registry enforces this at licence issuance rather than at execution.

Aggregate and derived workloads are the least constrained and the most valuable per unit of compute, since the sensitivity of the underlying material has already been reduced. These workloads can run wherever cost is lowest.

An increasing number of jurisdictions have enacted or proposed data localisation requirements affecting the processing of personal and sector-sensitive data, and several have created incentive structures for domestic compute capacity as a component of AI policy.

The mesh's commercial position is that it prices residency-constrained capacity at a premium while running unconstrained workloads at commodity rates, and that the margin between the two subsidises the Layer II price point. That cross-subsidy is the reason the compute layer is inside the protocol rather than procured from a hyperscaler, and it is also the reason the mesh is not designed to compete at frontier training scale.

6.2 Mesh Topology and Node Classes

Capacity is aggregated across three node classes with distinct economic roles.

Primary metropolitan clusters carry high-throughput inference and are positioned where network latency to the Anchor's densest enterprise concentration is low. They provide the baseline capacity for the majority of tasks and are intended to




reach cost parity with commodity regional inference.

Secondary regional clusters exist to satisfy residency boundaries that metropolitan clusters do not, and to serve contributor populations distant from the primary concentration. Their utilisation profile is lower and their unit cost higher, which is acceptable because the alternative for the demand they serve is a cross-boundary route that is not legally available.

Edge aggregation nodes provide low-latency pre-processing, data normalisation and residency attestation collection. They do not run frontier workloads. Their function is to reduce the volume of raw material that must traverse a boundary before the residency question is settled, which lowers both cost and compliance exposure.

Node operators post bonds sized against declared throughput rather than against reputation, since the failure mode that matters is an overclaimed capacity declaration. Residency attestation is a bonded claim as well: a node asserting that it operates within a boundary it does not occupy faces slashing at a severity level that makes the assertion not worth making falsely.

Interoperability with state compute programmes follows a narrow path. Where a programme has built capacity that is underutilised, the mesh can route compliant workloads to it under contract, which is attractive to both sides because the programme acquires utilisation and the mesh acquires capacity without capital expenditure.



6.3 Cost Structure Against Centralized Baselines

Cost comparison against centrally priced inference has to be conducted honestly, because the mesh does not win on every axis.




Aggregated regional capacity reaches a lower asymptotic cost per unit at moderate utilisation than centrally priced inference, primarily because the capital cost of the underlying hardware has already been borne by the node operators and the mesh pays only for delivered work plus a margin. That advantage narrows and then inverts at high utilisation, where the coordination overhead of distributed scheduling and the reliability variance across heterogeneous nodes outweigh the capital cost difference. The mesh therefore has a utilisation band in which it is clearly cheaper, and it should be positioned within that band rather than sold as universally cheaper.

Centralised baselines also carry surcharges that the mesh avoids. Cross-border data transfer fees apply to workloads that route outside a boundary. Compliance overhead attaches to any provider serving regulated demand and is priced into centralised rates whether or not a given customer requires it. The mesh's advantage in the residency-constrained segment is therefore larger than the raw compute differential suggests, and its advantage in the unconstrained commodity segment is smaller.

The gap between the two cost curves is the subsidy available to Layer II. That area is the figure's actual subject: it funds the difference between what agent delivery costs and what an enterprise below the formalisation threshold can pay. Where the subsidy is exhausted, either the enterprise price must rise or the capability ladder must stop at a lower level for that vertical, and both outcomes are matters of operating policy rather than protocol design.

The mesh reaches its cost advantage band at the utilisation levels the Phase II envelope produces. If utilisation runs above that band, the correct response is to route unconstrained workloads to external capacity while reserving mesh capacity for residency-constrained work, which preserves the subsidy at the cost of ceding the commodity segment.





7.0 The TLVY Token: Multi-Layer Utility and Value Capture

7.1 Design Axiom: No Utility Without Binding

A protocol token earns its place by doing something a contract, a database or an invoice cannot do at acceptable cost. Where a substitute exists, the substitute should be used, and the token claim should be withdrawn rather than preserved for symmetry.

This chapter therefore does two things in sequence. Sections 7.2 and 7.3 describe the mechanisms TLVY operates. Section 7.4 tests each of them against the substitute and reports which survive. The order matters: a necessity argument is only credible when written after the mechanism it interrogates, not before.

Three constraints govern every design choice here.

TLVY is never required to access a service. Enterprises pay for agent execution in local currency through a gateway. Licence counterparties pay in the currency of their contract. Contributors may elect settlement in TLVY or through a gateway conversion. Any design in which a user must acquire a token before receiving value would be rejected on usability grounds before regulatory ones.

TLVY carries no claim on protocol revenue. There is no share of licence income, no dividend, no redemption right, and no arrangement whose economic effect is equivalent even if its form differs. The waterfall in 4.4 distributes to contributors and validators as service compensation, not to holders as a class.

Every bonded position is forfeitable under stated conditions. A bond that cannot be lost is a deposit, and deposits do not secure anything.

7.2 The Five Utility Loops

Five distinct circulations involve TLVY, and they share participants without being parallel. The overlap is the design; treating them as separate functions would misrepresent how supply behaves.

The contribution loop. Verified work converts to TLVY credit at settlement. Contributors may hold, may bond into validator or node positions, or may convert to local currency through a gateway. The loop's significance is directional: value enters the token from outside the system through licence fees and agent task fees, then exits to contributors. Nothing in this loop generates TLVY; it distributes what the protocol has already earned or released.




The verification loop. Adjudicator and node operators post bonds to acquire adjudication and capacity rights, receive fees for service, and forfeit bond on adjudicated misconduct. The loop is a risk market. Its health is measurable in the ratio of fees earned to bond lost, and that ratio determines who participates.

The consumption loop. Agent operators pay task fees. A portion of the operator's payment covers inference and orchestration cost. The protocol's share is converted to TLVY at the gateway and a fixed percentage of the converted amount is burned, with the remainder held in treasury. Enterprise-facing pricing is denominated in local currency throughout, so the burn quantity varies with the TLVY market price while the enterprise's bill does not.

The licensing loop. Licence fees arrive in fiat, are converted for distribution, and route through the waterfall in 4.4. A share is burned at conversion. This loop is the most commercially significant and the least frequent per unit of activity — licence contracts are lumpy compared with task fees and carry much larger values.

The governance loop. Stake weights voting on the parameter set described in Chapter 10. Bonded stake time-locks for periods matched to the decision's reversibility, with constitutional matters carrying longer locks than operational ones.

The loops interlock at three shared points. Contributor settlement draws on the same TLVY pool that the licensing loop replenishes. Bonding removes TLVY from circulation and is replenished when validators exit and bonds unlock. The burn step reduces supply on both the consumption and licensing paths, which means the deflationary behaviour of the system is a function of two revenue lines rather than one.






7.3 Sink, Source, and Bond: The Ledger Mechanics

TLVY supply moves through four source channels and four sink channels, and the distinction between a sink and a lock matters more than it first appears.

Sources are cold-start contributor emissions, validator and node incentive reserves released on schedule, treasury releases authorised through governance, and ecosystem grants. The first is inflating by design and time-limited. The second is inflating and parametric. The third and fourth are reallocations rather than new supply at the margin, though they draw on balances that were themselves emitted.



Sinks are bonded stake, protocol fee burn, licence-fee conversion burn, and treasury absorption. Treasury absorption is the weakest of the four, since

absorbed TLVY can be released again and a reader should treat it as deferred supply rather than removed supply. Burn is the only sink that is irreversible, and bonded stake is the only one that is reversible but conditionally — a bond unlocks on schedule unless it has been slashed, which converts a probabilistic share of locked supply into removed supply.

Staking serves as an operational performance bond, not an interest-bearing instrument. Network distributions to node operators and Adjudicators represent dynamic service-fee settlements based on verified task volume and slashing risk, strictly maintaining a work-to-earn equilibrium rather than passive capital accrual. A validator who performs no adjudication earns zero fees.



Bond sizing scales with declared throughput for compute nodes, with adjudication volume for validators, and inversely with recent accuracy for both. The governing principle is that bond size should approximate the maximum value a participant could extract by misreporting their own conduct, which produces genuinely different figures across actor classes and across languages — a validator competent in a low-resource language adjudicates lower volumes and therefore carries a smaller absolute bond.

Slashed bond routes to one of three destinations. Restitution to an identified harmed party applies where a specific counterparty suffered loss. Burn applies where the misrepresentation has no identifiable victim, as with fabricated adjudication. Treasury applies to procedural violations without misrepresentation. Splitting the destinations this way avoids the failure mode of a single penalty pool, where every participant treats slashing as a shared operating cost rather than a transfer.

Bonding and Slashing Architecture					
Actor classes × Slashable offence classes × Slash destination × Severity weighting per cell - Empty cells show where slashing does not apply					
Actor Class	False Validation	Duplicate Submission	Capability Overstatement	Service-Level Breach	Data-Scope Violation
Adjudicators	High severity Wrong reject or false approve → Harmed party	High severity Fabricated adjudication → Burn	N/A	Medium Latency breach → Treasury	High severity Consent-scope mismatch → Harmed party
Compute Nodes	N/A	N/A	High severity Throughput overclaim → Burn	Medium Uptime SLA miss → Treasury	High severity Residency attestation false → Burn
Agent Operators	N/A	Low severity Task duplication → Treasury	Medium Template capability overstated → Treasury	Medium Accuracy SLA miss → Treasury	N/A
Gateway Partners	N/A	N/A	N/A	Medium Settlement delay → Treasury	High severity Screening bypass → Burn
Dispute Arbitrators	High severity Incorrect resolution → Harmed party	N/A	N/A	Low severity Latency breach → Treasury	N/A
High severity — restitution or burn Medium severity — treasury Low severity — treasury N/A — slashing does not apply					
DESIGN PRINCIPLE Matrix must show at least one empty cell class honestly — where slashing does not apply. Three-way split of destinations prevents treating all penalties as a shared operating cost. High-severity offences route to harmed party (restitution) or burn (no identifiable victim). Medium and low-severity procedural violations route to treasury.					
Calibrated discipline rather than punitive maximalism. Empty cells (N/A) show this is a designed system, not blanket enforcement.					

7.4 The Token Utility Necessity Test

The test is simple. For each function TLVY performs, ask what would happen to the protocol if TLVY did not exist and the function were performed by a database, an invoice and a contract. Where the substitute works, TLVY should not be paid for performing the function.



Function	Substitution without TLVY	Why the substitute fails
Contribution settlement across contributors with no prior contractual relationship	Manual payouts from a treasury account via banking rails	Cross-border micro-payout cost exceeds the value of the smallest contribution class, and provenance cannot be established from a payment record
Verification bonding	Reputational scoring alone	Reputation cannot be confiscated; false validation then carries no cost
Licence revenue distribution to a contributor pool of unknown and changing size	Centralised revenue-share arrangement	Requires identity resolution and payment corridors per contributor jurisdiction, which excludes informal contributors by construction
Supply discipline tied to utilisation	Discretionary buyback programme	Not auditable, not automatic, and dependent on a decision that can be reversed
Constituency for language admission and revocation policy	Foundation board appointment	No legitimate mechanism exists to weight the people whose material is being governed, and the resulting body becomes a political object

Five functions survive. That is a shorter list than most token documents present, and three functions that typically appear on such lists are absent here by choice.

Enterprise data licensing does not need TLVY. Licences are contracts between legal entities with identity verification, dispute resolution and reporting obligations. Settling them through a token would add settlement risk to a transaction that already has legal recourse, and would place a token inside the revenue line of the highest-margin counterparties.

Gateway settlement does not need TLVY. Regulated intermediaries already operate in fiat, and inserting a token into that path adds a conversion step that increases both cost and compliance complexity. TLVY's role in settlement is limited to the conversion step at which the protocol's own fee is realised, which is a treasury operation rather than a customer-facing one.

Agent subscription billing does not need TLVY. Enterprises pay in local currency. Tokenising the billing layer would require every customer to hold a volatile asset in inventory, which no operator in this category has the capacity to manage.



The design claim that follows is deliberately narrow. TLVY governs three things that cannot be handled by a database and an invoice — bonded security among parties with no legal relationship, settlement across payment-rail discontinuities at values below banking cost, and the constituency for language governance — and it is burned on the utilisation of two revenue lines rather than accruing a claim on them. A reader who concludes the token's scope is smaller than the protocol's scope has read the section correctly.

7.5 Value Accrual and Explicitly Rejected Anti-Patterns

TLVY's economic position depends on the interaction of three quantities: the rate at which revenue lines generate burn, the rate at which emissions release new supply, and the rate at which bonding removes supply from float without removing it from existence. The first is driven by commercial adoption, the second is scheduled and declining, the third is a function of network scale and can move quickly in either direction.

What TLVY holders hold is participation in a governance and security system whose parameter set determines how a licensing business operates. That is a real position with an economic value attached to the decisions it controls, and it is not a claim on cash flow. The distinction is worth stating without hedging because documents that blur it create liability for everyone involved.

Several mechanisms were considered and rejected, and the reasons are recorded because a reviewer will otherwise assume they were overlooked.

Buyback funded from treasury at management discretion was rejected on auditability grounds. The burn step occurs at a defined point in a published pipeline, and its size is a function of utilisation. A discretionary buyback is a promise about future behaviour.

Multiplied staking tiers offering higher rates for longer locks were rejected because they convert a performance bond into a yield product and attract capital that does not intend to adjudicate.

Bonded collateral in a second asset was rejected because it introduces correlation risk into the security layer at precisely the moment the second asset is most likely to fall.

Fee discounts denominated in TLVY for enterprise customers were rejected because they would pull the token into a customer-facing pricing decision, which is the boundary that the necessity test in 7.4 exists to maintain.



8.0 Supply Architecture: Distribution, Release, and Deflation

8.1 Total Supply and Issuance Discipline

Total supply is fixed at 10,000,000,000 TLVY. No mechanism exists to increase it. Cold-start emissions are drawn from a pre-allocated pool rather than minted, which means the emission schedule is a release schedule and the maximum supply is knowable at issuance.

That choice has a cost. A protocol that mints against contribution volume can scale rewards without bound; one that draws from a fixed pool cannot, and the pool's exhaustion is a hard constraint on the emission-funded phase. The decision to fix supply was made because the alternative introduces a governance question — how much to mint and when — that has no defensible answer, and because a fixed pool forces the phase gates in Chapter 11 to be real rather than extendable by argument.

Issuance discipline operates through three constraints.

Cold-start emissions decline on a published curve and terminate at a defined pool exhaustion point rather than continuing at a maintained rate. Validator and node incentive reserves release against service delivery, which is to say against evidence rather than against elapsed time. Treasury releases require authorisation through the governance path, and the treasury's total balance is a visible quantity rather than an implied one.

8.2 Allocation and Release Architecture

Ten categories hold the supply, and the reasoning behind each allocation is stated in terms of what the category has to fund rather than what it represents.


Category	Share	Mandate
Contributor rewards and cold-start emissions	28%	Funds verified contribution settlement until revenue covers it
Ecosystem and language development fund	14%	Language registry expansion, tooling, contributor onboarding in new varieties
Validator and node incentive reserve	12%	Bond subsidies and fee top-ups during the period before adjudication volume is self-sustaining
Core contributors	14%	Development team, four-year vest with a one-year cliff
Protocol treasury and operations	10%	Audit, legal, compliance, operating costs across the gated phases



Category	Share	Mandate
Liquidity formation reserve	8%	Draws against the tiers described in 9.2
Strategic partners and licensee onboarding	6%	Integration costs for licence counterparties and gateway partners
Public distribution	4%	Public liquidity bootstrap allocation
Advisors and subject-matter panels	2%	Domain and language expertise, subject to the same vesting as core contributors
Security reserve and contingency	2%	Held against exploit response, legal contingency and audit remediation

The contributor allocation is the largest single category, and it is scheduled to decline faster than any other. Cold-start emissions run at a high rate in the first period and fall on a curve that terminates rather than tapering indefinitely, because a contributor base that remains dependent on emissions has not produced a licensing business. The gate criteria in Phase II measure the ratio of licence-funded contributor settlement to emission-funded settlement directly, and a failure to move that ratio ends the emission programme rather than extending it.

Core contributor and advisor allocations carry a four-year vest with a one-year cliff. Validator and node reserves release against service delivery records. Liquidity formation reserve draws are event-driven and disclosed at draw rather than scheduled in advance, which is the only arrangement consistent with the flexible activation sequence described in 9.2.





8.3 Emissions Versus Burn: The Equilibrium Model

Emission volume is scheduled and therefore predictable within narrow bounds. Burn volume is episodic and keyed to utilisation across two revenue lines, which makes it the variable that determines whether supply contracts.


Three features of the burn mechanism shape the equilibrium.

Burn attaches at conversion rather than at collection. A licence fee is contracted in fiat, and the burn occurs when the protocol's share is converted for distribution. This means burn quantity moves with the TLVY market price while the underlying transaction does not, which produces an inverse relationship between price and burn rate that partially dampens both directions of movement.

The conversion ratio is a governance parameter with a published band. A higher ratio increases burn at the cost of treasury inflow, and the band constrains how far the ratio can move, so that a future governance majority cannot convert the mechanism into a near-total burn that starves operations.

Burn is asymmetric across the two revenue lines. Task fees are frequent and small; licence fees are infrequent and large. At low adoption, task-fee burn is a steady low-level trace against a much larger emission rate. As licensing scales, the episodic burns dominate the aggregate. The crossover in the model therefore arrives in two stages rather than one: first where licence burn exceeds emissions on an annualised basis, then where combined burn exceeds emissions on a trailing basis. Chapter 11's Phase III gate measures the first crossover, since it is the earlier and the more informative.

The utilisation level required to reach the first crossover is achievable within the Phase III envelope given the licence tariff bands in 4.4 and the modelled licence contract mix. The assumption is load-bearing and the falsification test is direct: if realised licence contract values fall below the modelled band by more than a stated margin, the crossover does not arrive inside the phase and the correct response is a narrower licence scope at a higher tariff rather than an extended emission programme.





9.0 Issuance Framework

9.1 *The Chain-Agnostic Canonical Supply Model*

TLVY is not issued on a chain. It is issued by a registry, and the registry deploys representations of itself to whatever execution environments serve the network's needs at the moment of a given decision.

That inversion is the substance of the Chain-Agnostic Canonical Supply Model. A single Canonical Supply Registry holds the authoritative record of total supply and the sole authority to authorise issuance on any environment. Deployment environments hold representations whose aggregate must reconcile against the registry at all times, and no environment can independently create supply in any form — no mint function, no rebase, no reissuance after burn, no wrapped derivative that expands the base.

The commercial reason for the arrangement is that the criteria which determine a good issuance environment in 2026 are not the criteria that will determine it in 2029. Three of them are currently in tension with each other.

Network liquidity determines how much depth a listing can attract and how quickly price discovery stabilises. Settlement cost determines the economics of the task-fee and licence-fee burn at the point of conversion. End-user access cost determines whether a participant in the Anchor archetype can interact with bonded positions without paying a meaningful share of their stake in transaction fees. A chain that scores well on liquidity may score poorly on the third criterion, and a chain that scores well on the third may not yet carry institutional depth.

Under the model, the selection of a primary issuance environment is a live decision evaluated against those criteria rather than a commitment made in a document. The candidate set is evaluated on a scoring rubric disclosed at the time of selection, and the outcome is announced with the reasoning rather than after it.

9.2 *Dynamic Multi-Tier Liquidity Formation*

Public distribution and capital formation proceed across four strata that may be activated in any sequence, in any combination, and at weights the protocol determines at the time of each activation rather than in advance.

Public liquidity bootstrap covers the formation of initial trading depth, whether on decentralised venues, through centralised listing arrangements, or through a combination of the two. It exists to establish a reference price with sufficient depth to absorb ordinary flow without material slippage.



Ecosystem contribution rounds cover allocation to parties delivering integration, tooling, language expansion and localised onboarding. Participation is against delivered work with milestone release, and the terms are structured as compensation for contribution rather than as investment.

Compliant institutional channels cover participation by entities subject to a regulated supervisory regime, whether through licensed intermediaries, structured vehicles, or facilities designed to meet jurisdictional investment-eligibility requirements. Structure and jurisdiction for any given activation are determined by the counterparty's own regulatory position rather than imposed generally.

Strategic market structure provisioning covers arrangements that improve the market's capacity to price and settle TLVY — market-making agreements, liquidity incentives tied to executed depth rather than to volume alone, and venue relationships. Where such arrangements are entered, their terms, duration and applicable conduct standards are disclosed.

Stratum	Participant profile	Lock character	Pricing formation
Public liquidity bootstrap	Unrestricted	None	Market discovery
Ecosystem contribution	Delivery counterparties	Milestone-vested	Negotiated against delivered scope
Compliant institutional	Supervised entities	Jurisdiction-determined	Structured against eligibility terms
Strategic market structure	Market operators	Performance-linked	Contractual against executed metrics


Sequencing, weighting and counterparty selection are determined under the discretion that disclosure obligations permit, which is a broader statement than a published calendar and a narrower one than an unbounded mandate. The specific constraint that applies at issuance is not the schedule but the disclosure: every activation carries a published account of which strata were used, what proportion of the liquidity formation reserve they drew, and what terms governed the release of any allocation.

Two design restrictions hold across all four strata. No stratum may be activated in a manner that would place the protocol's own distributions ahead of the public reference price in the settlement queue. And no stratum may be structured so that its participants receive supply at terms that a public participant could not have obtained had the terms been disclosed beforehand, except where the terms compensate for a verifiable service rather than for capital.



9.3 Canonicity and Supply-Integrity Controls

Reconciliation runs on a fixed cadence and on demand. The registry publishes total supply, aggregate authorised issuance, cumulative burn, and the balance held in each bonded position class. A discrepancy between the sum of environment-level representations and the registry figure halts new issuance authorisation until resolved, which is a blunt control and the correct one for the failure mode it addresses.




Bridge paths between deployment environments carry supply rather than creating it. A bridging operation locks a representation in the origin environment before authorising release in the destination, and the locked quantity is visible in the registry's reconciliation. Where an environment's bridge contract cannot provide this guarantee, that environment does not receive representation.

Burned supply is permanently retired at the registry. An environment that has burned TLVY holds a record of the burn, not a burn credit. Reissuance to replace burned supply is prohibited without a governance action that alters total supply, which is a constitutional matter under 10.2.

Bonded positions are recorded against the registry rather than against a single environment, which means a validator's bond cannot be double-counted across representations and cannot be evaded by exiting an environment. Slashing executes against the registry record, and unlock timing follows the registry.



The controls are heavier than a single-environment deployment requires. They exist because the model's advantage — the ability to move without reissuing — is also its principal risk, and the risk is concentrated in exactly the place where a reader of this document cannot independently verify anything: the bridge contracts that connect environments it may not know. Publishing the reconciliation cadence and the halt rule is the only available substitute for inspection.





10.0 Governance, Compliance, and Data Sovereignty

10.1 Governance Philosophy: What Token Holders Must Not Decide

Governance design is more about the decisions withheld than the decisions granted.

Three categories of decision sit permanently outside token-weighted voting. Clinical and safety-adjacent determinations do not arise in this protocol, but the equivalent category does: the quality-floor standards that determine whether a corpus is fit for licensing into regulated use. A body that can lower a quality floor to admit more supply can degrade the material that the protocol's entire commercial premise depends on, and the incentive to do so is proportional to the size of the contributor pool holding voting weight.

The second withheld category is adjudication outcomes in individual cases. Adjudicator panels decide disputes on their evidence. Token holders may set the rules under which panels operate, the severity weights applied to offence classes, and the appeal thresholds. They may not reverse a specific finding, and the absence of such a mechanism is a feature rather than a gap — appeals require a process, not a vote.

The third withheld category covers the protocol's obligations to counterparties with existing contracts. Licence terms already executed, gateway arrangements already in force, and disclosure commitments already published do not become revisable because a majority of holders would prefer otherwise. Governance can change what the protocol does next. It cannot reach backward into an instrument it has signed.

The reasoning behind the three prohibitions is not that holders would act badly. It is that the protocol's commercial value rests on counterparties being able to contract without pricing in the risk that a future vote changes the terms, and that risk is unpriced in a system where no such constraint exists.

10.2 The Three Chambers and the Unamendable Column

Three bodies hold decision authority, and their composition rules differ because the knowledge required for their decisions differs.

The **Holders' Assembly** comprises all staked participants weighted by stake and by lock duration. It sets emission parameters, treasury allocations within published bands, licence tariff bands by class, and admissions to the language registry. It also holds the authority to amend the framework itself, subject to the constraints below.




The **Adjudication Council** comprises bonded validators weighted by adjudication accuracy and volume rather than by stake alone. It sets quality floors by class, severity weights by offence type, and appeal thresholds. Inflating stake does not purchase weight here, which is deliberate: the body that determines whether material is fit for licensing should not be capturable by capital alone.

The **Technical Stewardship Committee** comprises engineering and security contributors appointed by the Assembly on the Council's nomination. It holds authority over protocol upgrades affecting the execution environment, the bridge controls in 9.3, and emergency pause. Seats are term-limited and no member may simultaneously hold a Council seat.

The **Unamendable Column** lists decisions that no vote of any body may alter, and it is populated. Whatever the Assembly's composition at any point, it cannot authorise supply above the fixed maximum, cannot create a redemption right or a revenue distribution, cannot remove the bonded requirement from any actor class, cannot alter the burn destination rules in 7.3, and cannot extend the emission schedule beyond the pool exhaustion point.

The supermajority required to amend the Unamendable Column is higher than the threshold for any other decision, and the procedure carries a requirement that the amendment be disclosed with a fixed interval between passage and effect. The interval is not a safeguard against bad faith. It is a safeguard against a decision made in a period of unusual market conditions being executed before reversal is possible.





10.3 Data Protection Architecture: On-Chain and Never-On-Chain

The registry records claims about material. It does not record the material, and the boundary between those two things is enforced structurally rather than by policy.

What exists on-chain: consent scope hashes, verification attestations, quality tier assignments, jurisdiction of processing allowed, retention limits, licence issuance records, revocation records, and the supply registry entries. All of these are small, non-reversible, and meaningful in aggregate.

What never exists on-chain: raw corpora, model weights, transcript content, speaker recordings, identity documents, and anything from which an individual contributor could be reconstructed. Consent scopes are recorded as hashes of a structured attestation held off-chain, so the on-chain record establishes that a scope exists and what class it belongs to without disclosing its terms to a public reader.

Erasure operates on the off-chain custody layer and propagates forward. A consent withdrawal removes the corpus from active storage, blocks new licence issuance, prevents registration of any derived asset, and records the revocation on-chain. What it cannot do is reach weights already trained, and Section 4.3 states this limitation rather than concealing it behind a technical description of what the mechanism does achieve.

Access control runs on a role basis with jurisdictional refusal rights. A gateway partner in one jurisdiction may be configured to refuse processing requests originating from a jurisdiction it is not authorised to serve, and that refusal executes at the gateway rather than at a central policy layer. The design consequence is that the protocol's compliance posture is partly delegated to intermediaries with direct regulatory exposure, which is efficient and introduces counterparty dependency that the risk register in 11.3 records.

10.4 The Settlement Gateway Layer and Multi-Gateway Custody Settlement

Protocol revenue accrues as a claim on fees paid in local currencies, and those fees are collected by regulated intermediaries operating under licences the protocol does not hold.


The Settlement Gateway Layer is the abstraction through which those claims reach the registry and through which distributions leave it. It performs four functions and holds no customer funds at any point.



It receives settlement instructions from licence counterparties and agent operators through whichever regulated channel is lawful in the relevant jurisdiction. It converts fiat fee proceeds to TLVY only at the point where the protocol's own share is realised – never for counterparty-facing transactions, which settle in local currency throughout. It executes screening against applicable sanctions and wallet-level requirements before any distribution instruction is released. And it reports to the registry in a format that reconciles against issued licences and recorded tasks.

The 4.4 waterfall executes across the gateway boundary rather than inside it. Gross fees enter, gateway and settlement costs are deducted at settlement, and the protocol's share converts at the disclosed conversion ratio with the burn applied to the converted quantity. A counterparty's bill is denominated in its own currency from quotation to settlement, and the token-level event is invisible to that counterparty – which is the arrangement that keeps the necessity test in 7.4 honest, since no customer is ever asked to price token exposure into a service contract.

Multi-gateway operation is the default rather than the exception. A single gateway concentrates operational risk and creates a jurisdictional single point of failure, so the design provides for several operating in parallel with per-jurisdiction configuration, per-jurisdiction refusal rights, and independent reconciliation against the registry. Each gateway's published conduct standards, licence status and reporting cadence are disclosed, and their failure to reconcile is a halting condition under 9.3.





11.0 Roadmap, Organization, and Disclosures

11.1 Phase-Gated Execution: Gates That Can Fail

Phases in this document carry no dates, and no quarter is named anywhere in the framework. A calendar is a promise about scheduling, and the protocol's progress depends on quantities it does not control: how many contributors a language community produces, how many validators can be attested in a variety that has never had them, how long a licence counterparty's procurement cycle runs. Fixing a date against those variables produces a document that ages badly and pressures the team toward announcing completion rather than measuring it.

Each phase is defined instead by gate criteria, each criterion is a quantity rather than an activity, and each has a stated consequence if it is missed. Three consequences exist. Re-scope narrows the phase's domain coverage and retries. Defer extends the phase without changing it, drawing on the treasury allocation set aside for that purpose. Halt ends the programme on the evidence available, and it is written into the framework as a legitimate outcome rather than a failure of resolve.

Phase I — Genesis Registry. The Commons is operational in a primary language and two adjacent ones. Mechanical verification resolves the linguistic material class without Adjudicator involvement. Median time from submission to settlement is measured and published. Contributor retention is measured after the cold-start emission rate tapers, which is the criterion that matters most and the one most likely to fail, since a contributor base that leaves when the emission rate falls was never contributing for the licence revenue. At least one executed licence contract, at any tariff class, with a counterparty that is not a related party.

A missed retention criterion triggers re-scope: fewer languages, deeper coverage in each. A missed licence criterion triggers defer, since the collection mechanism is not the constraint in that failure mode.

Phase II — Agent Commerce. Two verticals reach capability ladder level four, each measured on task volume and accuracy within its own domain rather than in aggregate. Contribution margin per enterprise is positive after channel share, on realised figures rather than modelled ones. The ratio of licence-funded to emission-funded contributor settlement is rising, and the gate requires a stated increase rather than an absolute level, because the absolute level at this stage is small for reasons that have nothing to do with the design. Channel partnerships are signed and producing onboarded enterprises, not merely announced.


A missed margin criterion triggers defer with the capability ladder held at level three for the affected verticals. A missed ratio criterion is the most serious outcome the framework anticipates short of a halt, because it indicates that the licensing business is not forming, and the response is a narrower licence scope at a higher tariff class.



Phase III – Mesh Consolidation. Residency-constrained capacity is live and serving regulated demand. Licence-fee burn exceeds emissions on an annualised basis, which is the first of the two crossovers described in 8.3. A second market within the Anchor archetype is served through a local operating arrangement, and the unit economics measured there are within a stated band of the Anchor's. Validator coverage is attested in at least five language varieties that began the programme with none.

Failure to reach the first crossover triggers defer once, then halt. The framework is explicit on this point: a protocol whose supply discipline never engages is a protocol with a permanent emission requirement, and continuing past that point spends contributor trust to buy time.

Phase IV – Federation. Independent operators run deployment environments under the Canonical Supply Registry's authority. External parties hold Adjudicator seats on attested competence rather than on appointment. Governance functions have moved from the stewardship entities to the three chambers described in Chapter 10, and the transition is verified by an external party rather than self-certified.





11.2 Expansion Matrix: Market Archetypes Beyond the Anchor

Expansion is described as a method applied to market types rather than a sequence of countries. The reason is partly discretion and mostly analytic: a claim stated as a property of an archetype can be tested in a market the protocol has not entered, whereas a claim stated about a country can only be verified by entering it.

Four structural criteria determine whether a market archetype is addressable, and they are the same criteria used in 2.4 with two additions that only matter after the Anchor is running. The two additions are the local cost of attestation, meaning whether competent validators exist in the target varieties or must be developed, and the presence of a channel structure through which enterprises below the formalisation threshold can be reached in groups.


The criteria behave differently in combination, which is what produces the sequencing.

Markets with high linguistic fragmentation and material micro-enterprise density but immature payment rails score well on collection value and poorly on Layer II billing. They are addressable through the Commons alone, with the agent layer deferred until a rail matures, and that is a legitimate configuration rather than a partial deployment.

Markets with mature rails and formal enterprise density score well on billing and poorly on collection value, since the workflows worth automating are already served by software with an incumbent relationship. These are the archetypes the protocol is least able to enter, and the matrix should say so.

The interesting band sits between them, and its defining property is that a fragmented contributor base and a mobile-primary enterprise base coexist in the same population. That property is not rare, and it is not evenly distributed.

One asymmetry governs sequencing across the matrix. The collection mechanism ports quickly: verification procedures, the taxonomy, the registry and the bonding architecture transfer to a new market with configuration rather than construction. The licence relationships do not port at all, because licence demand depends on a counterparty's own procurement rules and data-residency position. A market in which the Commons operates smoothly and no licence counterparty exists produces contributor obligations without revenue, and that configuration is the one the matrix is designed to prevent.



11.3 Entity Structure, Contributors, and Risk Factors

The protocol's legal architecture separates custody of the supply registry from the entities that operate revenue-generating activities, and it places the audit function outside both.

Entity	Function	Exposure
Development foundation	Holds protocol intellectual property; funds core development	Non-commercial; no customer relationship
Protocol stewardship entity	Holds the Canonical Supply Registry authority; executes governance outcomes	No operating revenue; acts on chamber instruction
Treasury vehicle	Holds reserves; disburses against authorised allocation	Assets disclosed on the published reconciliation cadence
Regional operating companies	Operate gateway relationships, agent deployment, contributor onboarding under contract	Carry commercial and regulatory exposure in their own jurisdictions
Independent verification function	Audits registry reconciliation, licence recording and burn execution	Structurally detached; contracted by the stewardship entity with a fixed term



Regional operating companies contract with the stewardship entity rather than owning protocol assets. That arrangement keeps the risk of a jurisdiction-specific regulatory action inside the entity that operates there. No operating company holds mint authority, registry access beyond its own reporting obligations, or the ability to alter licence records. Protocol assets and licensing income sit with the treasury and the development foundation, and no operating company retains licence revenue beyond its contracted service fee.

Core Architecture Working Groups

In alignment with decentralized governance principles and cross-border operational resilience, core development is led by specialized pseudonymous working groups coordinated through the Development Foundation:

- **Systems Architecture & Ledger Working Group (WG-Core):** Oversees the Canonical Supply Registry, cross-environment state proofs, and bridge contract security. Members contribute background in distributed fault-tolerant state-replication and high-concurrency clearing rails.
- **Computational Linguistics & Domain Evaluation WG (WG-Linguistics):** Directs the contribution taxonomy, inter-annotator consensus algorithms, and low-resource fine-tuning benchmarking. Focuses on regional language preservation and dialect corpus curation.
- **Cryptoeconomics & Mechanism Design WG (WG-Econ):** Manages dynamic bonding ratio models, parametric emission curves, and slashing schedule calibrations, ensuring mathematical equilibrium across the five utility loops.
- **Regulatory Engineering & Privacy WG (WG-Compliance):** Formulates consent-scope cryptography, residency attestation boundaries, and off-chain zero-knowledge compliance frameworks interfacing with regional gateway operators.


Specialized Domain Advisory Panels

Technical governance is augmented by independent subject-matter panels appointed on rotating, milestone-specific charters:

- *Linguistic Typology Panel:* Advises on annotator quality thresholds and phonetic dialect variations across target deployment clusters.
- *Institutional Data Governance Panel:* Advises on enterprise compliance standards, ISO-aligned data provenance, and regulated sector procurement interfaces.



Risk factors. The following are material and are stated with the residual risk that remains after the stated control.



Risk	Control	Residual
Contributor supply falls below the modelled volume in low-resource varieties	Tiered validation; validator development programme funded from treasury	The constraint is partly demographic; treasury funding changes its slope, not its floor
Validator concentration in a small number of attested varieties	Cross-language panel review; term limits on Adjudication Council seats	Concentration persists while attestation capacity is scarce
Licence demand does not form within the Phase II envelope	Narrower scope at higher tariff	The whole Layer I revenue model depends on this forming
Gateway counterparties fail, withdraw or become uneconomic	Multi-gateway default operation; per-jurisdiction configuration; independent reconciliation	Gateway dependency is structural for regulated settlement and cannot be engineered away
TLVY price volatility narrows the burn realised per conversion	Conversion ratio band; burn attaches at conversion rather than collection	A sustained price decline reduces burn quantity against fixed costs



Risk	Control	Residual
Bonding requirements become unattractive as the token depreciates	Bond sizing scales with volume, not with a fixed token quantity	In a deep drawdown, bond affordability and security strength move in opposite directions
Bridge or deployment environment failure severs supply reconciliation	Issuance halt on unreconciled discrepancy; no environment holds mint authority	A halted issuance does not recover an environment that has already failed
Regulatory reclassification of TLVY in a material jurisdiction	Utility framing limited to bonding, metering and governance; no revenue claim; no end-user holding requirement	Classification outcomes are not within the protocol's control
Emission dependence persists beyond Phase III	Fixed maximum supply; emission pool exhaustion is a hard constraint; halt defined as a legitimate outcome	The halt provision is only credible if it is used
Concentration of contributor allocation among high-volume early participants	Per-contributor accrual caps by quality tier; retention measured rather than assumed	Caps reduce the effectiveness of the reward for the contributors most capable of producing volume

Appendix A – Definitions and Terminology Index

These terms carry fixed meaning throughout. Where a term is used in ordinary language elsewhere, the protocol definition governs.

Contribution – A discrete unit of human work submitted to the Commons: linguistic material, domain annotation, evaluation judgement, translation, or structured local knowledge. Contribution is not equivalent to raw data upload.

Proof of Contribution – The consensus procedure by which a submitted contribution passes from provisional credit to settled status through stochastic



peer review, quorum verification and a fixed challenge window.

Adjudicator — The bonded validator class responsible for quality adjudication and dispute resolution. Adjudicators hold staked positions; their economic exposure is real and forfeitable.

Data Asset — A verified, deduplicated, consent-scoped corpus registered on-chain as a licensable unit. Carries a licence tariff band, a provenance record and an expiry or revocation condition.

Micro & Informal Enterprise (MIE) — An operating business below the formalisation threshold at which enterprise software is conventionally procured. Used throughout as an analytical category, not as a market segment narrative.

Canonical Supply — The single authoritative record of total TLVY supply and per-environment mint authority. No deployment environment holds independent issuance rights.

Settlement Gateway Layer — The abstraction through which protocol-level settlement instructions reach regulated local payment and custody intermediaries. The protocol does not settle in any jurisdiction directly.

Anchor Market — The initial deployment geography, characterised throughout by structural attributes rather than identity.