# Intelligent Internet Whitepaper

Emad Mostaque

Intelligent Internet

July 24, 2025

**Abstract**

This whitepaper introduces the Intelligent Internet (II), a novel protocol designed to address governance failures in AI development: the Great Race towards centralization and the Gradual Disempowerment of human agency. We propose a "Third Path" by building a Bitcoin for the Intelligence Age, gifting every human a sovereign AI (II-Agent), scaling intelligence through permissionless coordination via the Common-Ground protocol, and anchoring all knowledge on auditable, open-licensed datasets. The framework relies on Proof-of-Benefit, minting Foundation Coins only when verifiable societal benefit is performed. We outline a three-layer architecture (Foundation, Culture, Personal layers), define key actors and guiding principles (Openness, Verifiable Public Benefit, Credible Neutrality, Human + Agent Dignity, Sovereign Interoperability), and detail economic design, BFT finality, Anchor-Set provenance, Common-Ground orchestration, Universal AI access tiers, and progressive governance models. We provide a comprehensive security and threat model and discuss interoperability, ensuring the Intelligent Internet remains a transparent, auditable, and resilient public utility for the Intelligence Age.

# 1 Purpose and Vision

Rapid AI progress exposes two governance failures. The first is a **Great Race** toward centralized, monolithic AI controlled by a handful of actors. The second is a **Gradual Disempowerment** in which society trades away human agency for comfort. These are not destinies. These are failures of design. This whitepaper provides the engineering blueprint for a *Third Path*: the **Intelligent Internet**, a protocol intended to not merely solve technical problems but to offer a new social contract written in code.

## 1.1 Our Master Plan in Four Lines

1. **Build a Bitcoin for the Intelligence Age** by minting *Foundation Coins* only through **Proof-of-Benefit**.
2. **Gift every human a sovereign AI** in the form of an **II-Agent** bound to a non-custodial AI wallet.
3. **Scale that intelligence into permissionless coordination** through the **Common-Ground** protocol.
4. **Anchor everything on our shared inheritance of knowledge** with auditable, gold-standard, openly-licensed datasets whose full provenance is hashed on-chain.

Together, these four foundations create a self-reinforcing loop where societal benefit mints value. Value funds more capability, and capability remains transparently aligned with the common good. This is the "Bitcoin of the Intelligence Age".

*Symbioism* articulates the ethical North Star for these steps; the *Master Plan* states them succinctly; *this paper* turns them into protocol design.

## 1.2 The AI Paradox

The core of our challenge is a fundamental paradox: running large-language models is cheaper each quarter, yet access to trustworthy models is narrowing. This is not just a market inefficiency; it is a moral crisis driven by three interlocking trends:

- **Centralization**: Most inference traffic now passes through a handful of cloud providers, so a single outage or policy change can stall entire industries.
- **Opacity**: When model weights are closed, the public cannot inspect the data or reward signals that shape a system's behaviour; bias, theft, and sabotage stay hidden.
- **Fragility**: Tightly coupled SaaS chains break when a DNS blip, a regional block-list, or a legal injunction ripples through the stack.

Put simply, intelligence is getting cheaper, but its trustworthy supply is not.

## 1.3 Universal AI (UAI)

The Intelligent Internet responds with a guarantee: every person receives a baseline of open, verifiable intelligence. This is not a subscription service that rents your agency back to you; it is a system for sovereign intelligence you own and command.

- Each individual may create exactly one **II-Account**, a WebAuthn key whose hash is anchored on the ledger.
- That account comes with an **II-Agent**, a fully open-source reasoning agent that defaults to community models yet, under the user's key, can call specialist proprietary endpoints when required.

Every National Champion (the validator franchise for a jurisdiction) must honour a daily inference quota for the II-Agents it serves. A fully auditable, locally verifiable stack is therefore mandatory, not ideological.

### 1.4 Intelligence-Backed Capital

To power this new ecosystem, we redefine money. New tokens, **Foundation Coins (FC)** and **Culture Credits (CC)**, are minted only when validators attach a signed receipt that useful compute has been performed for the commons: data curated, models trained, or inference served. This is **Proof-of-Benefit**, not proof of waste.

Every coin is thus collateralized by a verifiable slice of cognition. FC follows Bitcoin's issuance curve yet halves twice as fast so that broad ownership can form before hardware gains widen inequality. In effect, FC is capital crystallized from intelligence.

### 1.5 Out-of-Scope Goals

The Intelligent Internet is a public utility, not a universal cure-all. It does not aim to build a global surveillance identity system, replace dedicated privacy-coins, outlaw proprietary AI, or settle every alignment debate. Its narrower mandate is clear: maintain an auditable, dependable layer of open intelligence, especially for the systems that bear society's greatest cognitive load.

## 2 Guiding Principles

The Intelligent Internet inherits its philosophical backbone from the Symbioism document but states each principle in operational terms so that engineers and regulators can test compliance without interpreting poetry.

### 2.1 Openness

Reference code, model weights, and datasets are released under OSI-approved licences. Anyone can rebuild the stack from source, reproduce hashes, and audit every line. This is more than a convenience; it is the only way a hospital, a bank, or a national archive can trust the software that now carries part of its cognitive load.

*Strategic Openness* follows: by lowering the cost of switching, an open stack deters zero-sum races toward closed, proprietary AGI.

### 2.2 Verifiable Public Benefit

The ledger mints FC only when validators provide cryptographic evidence that useful compute (dataset curation, training, inference) was performed for the commons. No proof, no reward. Scarcity is therefore tied directly to measurable public benefit.

### 2.3 Credible Neutrality

The protocol offers no hidden levers: validator admission rules are objective, economic policy is on-chain, and planned upgrades require a global, replayable vote with a 30-day grace period. Power is distributed, and when it must exist, transparent.

### 2.4 Human + Agent Dignity

Every citizen receives one II-Account and an II-Agent that runs under their key. Core functionality does not require KYC, and optional proofs (personhood, local compute, external APIs) grant only incremental privileges. Guardian Sentinels monitor code paths so that private data does not leak upstream.

### 2.5 Sovereign Interoperability

Nations and sub-cultures can extend the stack without forking its root. Culture layers levy fees in their own CC units and may enforce local data-residency rules, yet a light client anywhere in the world can still verify Foundation layer headers and prove ownership of FC.

These five principles serve as the constitution of the project: every technical detail that follows must trace its lineage back to at least one of them.

## 3 System Model

This chapter describes the actors, the three-layer architecture, and the trust boundaries the protocol assumes at genesis.

### 3.1 Actors

- **Citizens** generate a single II-Account, a WebAuthn key whose hash is anchored on the ledger.
- **II-Agents** are software counterparts to citizens. A top-level Partner agent orchestrates one or more Principals, which spawn specialized Associates to execute discrete work-modules. Agents are fully open-source, yet, under the user's key, they can call proprietary endpoints when a task demands it.
- **National Champions** are validator franchises that run large compute clusters, finalize Foundation layer blocks, and guarantee the daily UAI quota for citizens in their jurisdiction. The network launches with twelve Champions; the long-term goal is at least one per sovereign nation.
- **Guilds** curate domain datasets and benchmarks (medical, legal, linguistic) and can veto a dataset's inclusion in an Anchor-Set if it breaches licensing or privacy.
- **Guardian Lattice** is a set of advanced II-Agents operating in three roles. Sentinels monitor code paths and SLAs, Advisers raise alerts, Implementers can pause the network in an emergency.
- **Oracle Council** is a fifteen-seat body, takes over parameter tuning (benefit classes, SLA thresholds, job-class registry) once Phase Three begins.

### 3.2 Three-Layer Architecture

- **Foundation layer** (L0) is the canonical ledger. It records UTXOs, FC issuance, Anchor-Set roots, and identity hashes.
- **Culture layer** (L1) runs as sovereign roll-ups. Each Champion maintains a roll-up that settles in its own CC unit, enforces local data-residency rules, and forwards UAI queries to the Foundation layer when global state is required.
- **Personal layer** (L2) lives on user devices. Here II-Agents execute workflows, keep private context, and decide whether to call a local model, a Champion endpoint, or an external API.

*Implementation note*: the reference node software begins as a fork of **Bitcoin Core v25**, with amended BFT consensus. Existing Bitcoin tooling therefore needs only minimal adaptation. This will likely evolve as the network heads towards public release.

A light client can verify Foundation layer headers with Merkle proofs alone; no jurisdiction can mandate a proprietary full-node binary.

### 3.3 Trust and Adversary Assumptions

- **Honest-super-majority**: At least two-thirds of Champion stake behaves correctly at any moment.
- **Bounded network delay**: The Champion Interlink fabric provides sub-300 ms round-trip; temporary partitions heal within that bound.
- **Cryptographic soundness**: Hash functions and signature schemes are unbroken; quantum attacks are not practical before the planned Dilithium migration.
- **Rational economics**: Champions value unreduced stake and an active licence more than short-term equivocation or SLA evasion.

All consensus rules (block validity, slashing, parameter upgrades) rely only on these minimal assumptions. If any are violated, the system degrades gracefully to "pause and replay from last undisputed checkpoint," matching the resilience of battle-tested open blockchains.

## 4 Proof-of-Benefit (PoB)

Every new FC is minted only when a block carries verifiable evidence that the network has delivered a measurable public benefit. The first benefit class is **open-stack compute** (dataset curation, model training, or inference) but the framework allows additional classes to be authorized over time.

### 4.1 Recognized Benefit Classes at Genesis

- **Compute-Inference**: Serving end-user requests for models listed in the public registry.
- **Compute-Training**: Training or fine-tuning those models on auditable data.
- **Data-Curation**: Running the II-Commons pipeline: deduplication, contamination checks, vector indexing.
- **Agent-Orchestration**: Submitting execution traces that show an II-Agent completed a work-module under the $\geq$ three-module rule.

Future classes (energy-load balancing, zero-knowledge verification services, public-domain content generation) can be added by an Oracle-Council vote; no hard-fork is required.

### 4.2 PoB Receipt: Minimal Fields

Every unit of recognized benefit emits a receipt object:

- **work_root**: Merkle root committing to inputs, outputs, and artefacts.
- **benefit_class**: String such as compute_inference or data_curation.
- **eval_score**: A numeric score produced by the evaluation harness; each class has a threshold set by Oracle Council.
- **producer_id**: Public-key hash of the Champion (or delegated pool) that performed the work.
- **timestamp**: Unix seconds at completion.

The full schema lives in the *AI Methods & Stack* note (§5). Foundation layer consensus cares only that the fields exist and that eval_score clears the active threshold.

### 4.3 Block Validity Rule

A Foundation layer block is valid only when all conditions below hold:

1. $\geq$ two-thirds of the validator stake signs the block header.
2. The header embeds precisely one PoB receipt with an eval_score meeting the threshold for its benefit_class.
3. The same validator quorum signs the receipt hash.
4. Random availability sampling of work_root leaves by at least $m$ validators succeeds ($m$ TBD).

No receipt, or a receipt below threshold, means no new FC is minted and the block is rejected.

### 4.4 Issuance Epochs and Halving Schedule

- **Epoch length**: 105,000 valid blocks.
- **Halving**: The block subsidy halves after every two epochs. Because an epoch is block-count based, calendar duration self-adjusts if block time changes.
- **Fixed cap**: Total supply is hard-capped at 21 million FC.

### 4.5 Threshold Management

Initial thresholds for each benefit_class are set at genesis. Thereafter:

- The Oracle Council may raise, lower, or add thresholds by a two-thirds vote.
- Changes are announced on-chain and activate after the 30-day governance grace period defined in Chapter 10.

### 4.6 Fraud-Proof and Dispute Path

If a validator suspects a fraudulent receipt (fabricated outputs, forged work_root, or inflated eval_score) it submits a dispute transaction:

1. The disputing validator posts the receipt hash and claim.
2. The producing Champion has a fixed window to disclose artefacts.
3. Guardian Sentinels rerun the reference verifier.
4. If fraud is confirmed, the Champion is slashed under Chapter 6 rules and the disputed reward is burned.

Chapter 4 establishes one invariant: **no verifiable public benefit, no new coins.** All economic and security guarantees later in the paper rely on that link.

## 5 Economic Design

The Intelligent Internet turns verifiable benefit into currency and currency back into more benefit, forming a closed economic loop. Energy finances compute, compute produces open knowledge, open knowledge mints Foundation Coins, and Foundation Coins in turn finances the next round of compute and infrastructure.

### 5.1 Foundation Coins (FC)

- **Fixed cap**: Total supply can never exceed 21 million.
- **Issuance curve**: At genesis the block subsidy matches Bitcoin's 50 FC per valid block but the reward halves after every two issuance epochs instead of four. An epoch is 105,000 **valid** blocks, so if block time accelerates the calendar interval contracts automatically. Under these rules slightly more than 99 % of FC will be in circulation about twelve years after mainnet launch.
- **Backed by benefit**: Because a block is valid only when it embeds a qualifying Proof-of-Benefit receipt, every minted coin is collateralized by a measurable public good.
- **Unit of account for settlement**: All Foundation layer transaction fees are paid in FC; National-Node fees may be paid in FC or the relevant Culture Credits.

### 5.2 Culture Credits (CCs)

- **Purpose**: Culture Credits give each jurisdiction a monetary layer that can respect local policy (taxation, privacy, data-residency) without fragmenting the global ledger.
- **Mandatory fee**: Every National-Node transaction must carry a fee in the local CC; the initial rate may be symbolic and can be adjusted by national policy over time.
- **Reserve requirement**: Each CC must hold a floor of FC as backing collateral. The exact percentage is still open and will be set after public modelling and consultation.
- **Peg mechanics**: If a CC drifts below its reserve floor, its roll-up must raise fees, auction debt, or restrict issuance until the peg is restored.
- **Interoperability**: Edge wallets can swap CC – FC through on-chain automated market makers; a light-client can verify the peg without running a full National-Node.

### 5.3 Treasury

- **Revenue sources**: A programmable share of every FC block reward and Foundation layer fee flows into the Treasury. National-Nodes may contribute an additional share of CC fees.
- **Burn and grant**: Treasury logic will burn a fraction of incoming FC to offset issuance and allocate the remainder as grants for public datasets, model training, infrastructure upgrades, and open-source tooling. The exact burn–grant split is still to be set by governance before mainnet launch.
- **Governance**: During Phases 0–2 a multi-sig held by the founding Champions stewards Treasury outflows. At Phase 3 control passes to an on-chain process supervised by the Oracle Council and, where relevant, Guilds.
- **Transparency**: All Treasury transactions settle on the Foundation layer; the contract exposes a JSON view so that any wallet or regulator can audit inflows, burns, and grants in real time.

Together these three components: benefit-backed Foundation Coins, policy-scoped Culture Credits, and a transparent Treasury, close the economic flywheel that powers the Intelligent Internet: **power becomes compute, compute becomes open knowledge, open knowledge becomes currency, and currency finances more power and compute.**

## 6 BFT Finality and Proof-of-Benefit Attestations

The Foundation layer ledger finalizes transactions through a Byzantine-Fault-Tolerant vote that embeds a Proof-of-Benefit (PoB) receipt in every block. The design keeps Bitcoin Core's validated data structures

but replaces proof-of-work with one-round BFT voting plus benefit verification. A block that lacks a valid PoB receipt never reaches finality and mints no Foundation Coins.

### 6.1 Validator Set at Genesis

The network launches with twelve **National Champions**. Each Champion posts an economic stake. The amount $\mathbf{S}_{\mathrm{FC}}$ will be fixed before mainnet, along with a hardware-attestation report. The long-term aim is at least one Champion per sovereign nation; the admission process is open and rule-based.

### 6.2 Block Lifecycle

1. **Proposal**: A round-robin leader packages pending transactions and exactly one PoB receipt into a candidate block.
2. **Vote**: The header and receipt hash are multicast across the Champion Interlink Network. Validators sign if both pass stateless checks.
3. **Commit**: The block becomes final when signatures representing at least two-thirds of bonded stake arrive. There are no subsequent re-organizations.
4. **Broadcast**: The quorum certificate and receipt metadata propagate to National-Nodes and light clients.

Because the protocol inherits HotStuff safety proofs, correctness holds as long as less than one-third of stake colludes or is offline.

### 6.3 Champion Interlink Network

To sustain single-round finality, every Champion maintains direct connections to all others:

- End-to-end round-trip latency under 300 ms.
- Symmetric bandwidth of at least 10 Gbit/s on three independent physical paths.
- Signed heart-beats every second so Guardian Sentinels can measure uptime.

Violations are logged to the Sentinel channel and may trigger slashing when thresholds are exceeded.

### 6.4 Slashing and Licence Revocation

A Champion forfeits **N %** of its stake (parameter TBD) and loses its franchise licence if it:

- **Equivocates**: Signs two different blocks at the same height.
- **Misses**: More than **k** consecutive votes or proposals within a sliding window $\Delta$.
- **Undershoots SLA**: Submits PoB receipts whose evaluation scores fall below threshold **S** for their declared benefit class.

Once a licence is revoked, the validator may still operate on the Personal layer with an II-Agent, but it is ineligible for future FC rewards until it wins an open tender for the vacant slot.

### 6.5 Admission and Exit

- **Admission**: An aspiring Champion posts stake $\mathbf{S}_{\mathrm{FC}}$, publishes a hardware-attestation proof, and passes a one-week public objection period; final approval requires a Sentinel majority.
- **Exit**: A voluntary exit triggers a cool-down of $\tau$ blocks (value TBD). Stake unlocks only after all open dispute windows on that validator's receipts close.
- **Vacancies**: If a licence is revoked or a validator exits, an open tender admits a replacement. Selection is based on stake, technical readiness, and projected SLA compliance.

### 6.6 Safety, Liveness, and Energy Notes

Under the latency bound the protocol reaches finality in a single communication round. If a network partition violates the bound, honest validators refuse to sign conflicting headers; the worst outcome is a temporary stall until connectivity is restored.

Champions are encouraged, though not yet required, to self-attest that a majority of their power draw comes from renewable or stranded energy sources. The Oracle Council may formalize third-party audit requirements once reliable metering standards mature.

Chapter 6 defines how a block becomes irreversible, how public benefit is enforced at the consensus layer, and how the system removes misbehaving validators. All economic and governance guarantees that follow rely on these properties.

## 7 Anchor-Set Overview

The Anchor-Set mechanism gives the ledger a durable memory of the data and model artefacts that power Universal AI. Every dataset, index, model checkpoint, and II-Agent identity is reduced to a Merkle root and pinned to the Foundation layer chain. Anyone can audit provenance, verify integrity, and reproduce results without storing raw data on-chain.

### 7.1 From Raw Data to Anchor-Set Root

1. **Source collection**: Web, academic corpora, government releases, and Guild-contributed archives enter an ingestion queue.
2. **II-Commons pipeline**: The queue is deduplicated, de-contaminated against popular benchmarks, and embedded into vector indexes. Each stage logs hashes referenced later.
3. **Common Ground QA**: Human reviewers, working through II-Agents, spot-check samples, tag licence metadata, and red-team for disallowed content.
4. **Root calculation**: A Merkle tree covers raw shards, vector files, QA logs, and licence tags. Its root, the **Anchor-Set root**, is ready for on-chain commit.
5. **Commit**: Champions include the root in a special OP_ANCHORSET output.

For large models the process is similar: weights are sharded, hashed, and overall root is committed.

### 7.2 Agent Identity Roots (OP_IDCLAIM)

Every II-Agent has a long-term public key. Once per rotation epoch, at least annually, the agent publishes the SHA-256 hash of that key in an OP_IDCLAIM transaction. The same hash is relayed to the relevant National-Node roll-up, allowing local policy without a duplicate commitment.

If a key is compromised, the user rotates keys and posts a new hash; the old root remains for historical linkage, but only the latest root is considered active for telemetry and status upgrades.

### 7.3 Governance and Versioning

- **Sentinel + Champion sign-off (Phases 0–2)**: At launch, a new Anchor-Set root becomes canonical when a majority of Guardians and two-thirds of Champions include the same root in a window of 1,000 blocks.

- **Oracle Council escalation (Phase 3 onward)**: After decentralization, the fifteen-seat Oracle Council can approve or reject a root by a two-thirds vote. Rejections must cite a licence, privacy, or contamination breach and are recorded on-chain.
- **Version tags**: Roots carry semantic versions (vMajor.Minor.Patch). A Minor bump adds or removes shards; a Patch bump fixes metadata only.

### 7.4 Rotation and Expiry Policies

- **Dataset and model roots**: Must be refreshed at least once every twelve months or when a licence or contamination change is detected.
- **Identity roots**: Must rotate at least annually; sooner if keys are compromised.
- **Old roots**: Remain accessible for reproducibility but fall out of the "active set" used by PoB verification once superseded.

### 7.5 pubflag Logic

An Anchor-Set commit carries a single-bit **pubflag**:

- pubflag = 1 $\rightarrow$ raw data or weights are openly downloadable under their stated licence.
- pubflag = 0 $\rightarrow$ only a salted hash is on-chain; raw shards stay encrypted. Guardians still run deduplication and contamination proofs without revealing content.

The pubflag lets privacy-sensitive datasets join the benefit pool while satisfying legal constraints.

### 7.6 Light-Client Verification

Edge wallets and National-Node roll-ups can verify an Anchor-Set root by pulling:

1. The Foundation layer header that contains the OP_ANCHORSET.
2. A Merkle proof linking the root to the header's transaction merkle root.
3. The licence and timestamp in the root's metadata.

No full node is required; a phone can confirm provenance with kilobytes of data.

Anchor-Sets give the Intelligent Internet a tamper-evident history of its training data, models, and agents. They are the foundation for transparent audits, fine-grained governance, and reproducible science.

## 8 Common-Ground Orchestration

Every complex task on the Intelligent Internet is broken into small, auditable pieces and executed by a team of II-Agents following the **Common-Ground** protocol. The design borrows ideas from operating-system process scheduling and distributed build systems: plan first, parallelize aggressively, record everything.

### 8.1 The II-Agent Hierarchy

- **Partner (top-level) II-Agent**: Lives on the user's Edge device, interprets natural-language goals, keeps close communication with the user, drafts project plan and oversees execution.
- **Principal II-Agents**: Spun up by the Partner; orchestrates the execution team towards the project goal, such as *"summarize regulatory filings"* or *"fine-tune a domain model"*.
- **Associate II-Agents**: Short-lived workers that execute individual work-modules: web search, code generation, data cleaning, GPU training runs.

All agents are fully open-source. When a proprietary endpoint offers superior accuracy, the caller must supply its own API key so that proprietary terms never leak into the public licence chain.

### 8.2 Work-Module State Machine

A work-module moves through four states:

1. **pending**: Drafted by the Partner or Principal.
2. **running**: Assigned to an Associate and underway.
3. **pending_review**: Finished; awaits Principal verdict.
4. **completed / deprecated**: Accepted or discarded.

A project **must** begin with at least three modules, and agents **should** group independent modules into a single dispatch call so that Associates run them in parallel.

### 8.3 Telemetry and Audit Trail

Each agent call **must** emit a telemetry event that includes:

- **agent_id_root**: The latest OP_IDCLAIM hash for that agent.
- **module_id**: UUID of the work-module.
- **steps**: Token counts or GPU hours consumed.
- **timestamp**: Unix seconds.

Champions relay telemetry to the Guardian Lattice in real time. The default emission rate is one event per tool call; the Oracle Council may tighten or loosen the rate before public launch.

### 8.4 Opt-Out Flag for Sensitive Artefacts

If a project handles content that must remain private (e.g., patient records), the Partner agent can set an **opt-out flag** on the module. The flag prevents automatic Anchor-Set commits of intermediate data while still allowing a salted hash to satisfy PoB verification. Guardians confirm that the hash meets deduplication and contamination tests without exposing raw content.

### 8.5 Benefit Eligibility

- **Compute-Inference receipts**: Generated when Associates serve inference queries.
- **Data-Curation receipts**: Arise from deduplication or vector-index tasks.
- **Agent-Orchestration receipts**: Cover the coordination overhead itself (planning, scheduling, and verifying results) provided the project used at least three parallel modules.

All receipts carry sla_class = agent_orchestration and an evaluation score derived from module-completion latency, accuracy checks, and resource efficiency.

### 8.6 Failure Handling

If an Associate stalls or errors:

- The Principal marks the module pending again and rewrites its description.
- The Partner may spawn additional Associates or split the module further.
- A repeated-failure counter increments; exceeding the threshold lowers the eval_score for the final receipt, putting the Champion's stake at risk.

By enforcing structured planning, parallel execution, and exhaustive logging, Common-Ground turns loose collections of LLM calls into reproducible, auditable workflows. This is exactly what regulated industries require when they hand critical tasks to machine agents.

## 9 UAI Access and Status Tiers

Universal AI (UAI) turns the ledger's open models and datasets into practical, everyday intelligence for individuals. This chapter explains how a citizen obtains an II-Account, how daily access is guaranteed, and how greater capability can be unlocked.

### 9.1 II-Account Creation

Every person may generate one II-Account by creating a WebAuthn key-pair on their own device.

The public key's SHA-256 hash is posted to the ledger in an OP_IDCLAIM transaction and becomes the canonical identity for that citizen's II-Agent.

No compulsory KYC is required for this base step, and the hash contains no personal data.

If the key is ever lost or compromised, the holder can rotate to a new key: a fresh OP_IDCLAIM supersedes the old hash and preserves account continuity.

### 9.2 Baseline Daily Quota

Each National Champion must honour a published minimum of inference tokens per citizen per day.

Champions can satisfy the quota with their own GPUs, with cached model shards pushed to edge devices, or by load-sharing with neighbouring Champions.

When a Champion serves a request, it tags the corresponding telemetry record as benefit_class = compute_inference, making that work eligible for PoB minting when the next block finalizes.

### 9.3 Status Upgrades

Citizens who need more capacity than the baseline (or who wish to run specialized tasks) can raise their status tier in three independent ways:

- **Proof-of-Personhood (PoP)**
  - Privacy-preserving attestation, such as a zero-knowledge passport check, confirms the user is a unique human.
  - Champions reward PoP holders with a higher daily token cap and shorter queue times.
  - The proof itself is never stored on-chain.
- **Donated Local Compute**
  - The user allows their device to run jobs scheduled by a Guardian Sentinel. GPU hours are measured and signed locally.
  - Contributing compute earns priority access to new models and may even stream micro-rewards of Foundation Coins.
  - Only an aggregate resource log leaves the device.
- **External API Keys**
  - A user can attach their own licence key for a proprietary model (legal, medical, or otherwise).
  - The II-Agent calls the endpoint under the user's key and merges results with open-model output.
  - The key remains on the device and is never broadcast.

An II-Agent rolls these signals (PoP status, compute contribution, and external licences) into a single **status score** that it presents when requesting service from a Champion. Self-attestation alone is sufficient for baseline access; upgrades are fully optional.

### 9.4 Alignment Tailoring

Because each citizen controls exactly one II-Agent, the agent can learn personal context (reading level, accessibility needs, ethical preferences) without exposing data upstream. Alignment and capability scale with the status score, but core safety rules remain identical across all tiers.

### 9.5 Rate-Limit Enforcement and Abuse Mitigation

Champions enforce quotas at their API boundary. Excess requests receive a 429 response and, if appropriate, a link explaining upgrade options.

Sybil resistance relies on the economic cost of stake and compute: mass-created agents without PoP or contributed cycles remain at the baseline cap.

Telemetry records anonymized misuse flags, such as attempts to generate malicious code, and repeated abuse can lower an agent's status score after Guardian review.

UAI guarantees that every citizen enjoys a dependable baseline of open, auditable intelligence. Those who contribute additional value (identity proofs, spare compute, or proprietary licences) unlock proportionally greater capability, all without compromising privacy or credible neutrality.

## 10 Governance and Alignment

### 10.1 The Tripod of Justice

Symbioism frames legitimacy around three tests: **openness** (isolation is forbidden), **flow** (stagnation is prohibited), and **resilience** (transgression is self-defeating). Every rule in this chapter exists to keep that tripod level.

### 10.2 Guardian Lattice: II-Agents on Watch

The Guardian Lattice is staffed by specialized **II-Agents** running under three mandates:

- **Sentinels** stream runtime logs, meter energy attestations, and replay PoB receipts.
- **Advisers** digest Sentinel telemetry and publish human-readable risk reports.
- **Implementers** hold a time-locked pause key that can stop block production in an emergency.

Rotation is automatic: every six months a stake-weighted election re-deploys fresh Sentinel, Adviser, and Implementer agent images, signed and reproducible like any other open-source model.

### 10.3 Oracle Council: Common-Ground in Miniature

Once Phase Three begins, parameter governance (new benefit classes, SLA thresholds, job-class registry) moves to a fifteen-seat Oracle Council. Each Council seat is itself a **Partner-grade II-Agent** orchestrating a micro-team of Principal and Associate agents through the Common-Ground protocol:

- A parameter proposal arrives as a work-module plan.
- The Council-agent decomposes analysis into parallel Associate tasks: benchmark replay, licence scan, economic-impact simulation.
- Results roll up to the Council-agent, which publishes a deterministic vote.

Human overseers still sign the final transaction, but the heavy lifting (data gathering, simulation, and evaluation) is executed by auditable, open-weight agents. A proposal passes when at least ten Council votes concur, and the change activates after the thirty-day grace period.

### 10.4 Planned Upgrade Flow

1. A Champion, Guild, or Sentinel submits an on-chain parameter diff.
2. Sentinel agents simulate the change and publish a pass/fail report within seven days.
3. Oracle Council agents run their Common-Ground plans and cast votes.
4. If two-thirds approve, the network waits thirty days, then Champions embed the new parameter blob in the next block header.

### 10.5 Progressive Decentralization Timeline

- **Phase 0 – Spec Freeze**: Founders publish code; testnet launches.
- **Phase 1 – Public Testnet**: Guardian agents elected; live telemetry begins.
- **Phase 2 – Mainnet**: Twelve Champions secure the Foundation layer; Treasury under multisig.
- **Phase 3 – Oracle Council Hand-off**: Council agents assume parameter control; Treasury grants become on-chain proposals executed by Common-Ground.
- **Phase 4 – One Champion per Nation**: Validator set expands; Guardian and Council seats scale accordingly.

By grounding every monitoring and governance role in the same auditable II-Agent framework that serves end-users, the protocol keeps human discretion at the policy layer while pushing routine evaluation and enforcement down to transparent, reproducible code.

## 11 Security and Threat Model

This chapter catalogues the principal risks to the Intelligent Internet and the counter-measures built into the protocol, the hardware assumptions, and the governance process.

### 11.1 Adversary Categories

- **Network adversary** aims to censor or reorder blocks by partitioning the Champion Interlink or delaying messages beyond the 300 ms bound.
- **Validator adversary** seeks to double-spend or inflate supply by equivocating, submitting fake Proof-of-Benefit receipts, or bribing other validators.
- **Data-poison adversary** injects contaminated or malicious data into Anchor-Sets to degrade model quality or expose private information.
- **Model-poison adversary** attempts to insert back-doored weights or to compromise inference endpoints.
- **Cloud fragility centralized providers** may revoke API access or succumb to outages, leaving users stranded.

### 11.2 Protocol-Level Defences

- **One-round BFT finality** ensures that even with network delay an honest-super-majority cannot be tricked into signing conflicting histories.

- **PoB receipt signatures** anchor the mint path to measurable public benefit; falsifying a receipt exposes the validator to slashing through the dispute mechanism in Chapter 4.
- **Availability sampling** obliges validators to hold or at least locate the artefacts referenced by a receipt, deterring "hash-only" frauds.
- **Automatic slashing and licence revocation** remove malicious or unreliable Champions and re-tender the slot without manual intervention.

### 11.3 Data and Model Integrity

- **Anchor-Set Merkle roots** make any tampering with training data, vector indexes, or model weights obvious to downstream validators and auditors.
- **Guardian Sentinel replay** reproduces evaluation scores for receipts and flags divergence. If a new benefit class is approved, its verifier code must be deterministic and open-source.
- **Opt-out flag with salted hash** permits the use of sensitive datasets while still allowing deduplication and contamination proofs to run.

### 11.4 Sybil and Spam Resistance

- **One II-Account per citizen** anchored on-chain prevents unlimited account farming.
- **Baseline inference quota** is guaranteed yet capped; higher tiers require proofs-of-personhood or donated compute, raising the cost of Sybil spam.
- **Telemetry misuse tags** let Sentinels downgrade or throttle agents that repeatedly request disallowed content.

### 11.5 Energy and Sustainability Considerations

Champions self-attest that a majority of their power draw comes from renewable or stranded energy. Audit hashes for third-party reports are encouraged and may become mandatory once reliable metering standards mature. Because no proof-of-work exists, energy use scales with the social value of compute rather than with an arbitrary hash-rate race.

### 11.6 Post-Quantum Readiness

The ledger currently uses ECDSA-secp256k1 signatures for compatibility with Bitcoin-derived tooling. A migration path to **Dilithium** (or a successor PQ signature) is defined in the technical annex. Activation requires:

1. A client implementation that can verify both signature schemes.
2. A Sentinel simulation report showing throughput impact.
3. An Oracle Council vote followed by the standard 30-day grace period.

With these defences, the Intelligent Internet maintains integrity, availability, and auditability in the face of realistic adversaries, while preserving a clear migration path for emerging threats such as quantum cryptanalysis or large-scale cloud failures.

## 12 Interoperability and Ecosystem

The Intelligent Internet is designed to exchange value, data, and compute with existing blockchains and web services while safeguarding its own consensus and benefit-mint cycle.

### 12.1 Light-Client Verification

Any external chain or wallet can verify Foundation layer headers with only kilobytes of data:

1. Download the header that contains the block hash, quorum certificate, and PoB receipt hash.
2. Check the aggregated BFT signatures against the published validator set.
3. Verify the PoB receipt hash against the header's transaction Merkle root.

Because the flow mirrors Bitcoin SPV, most existing libraries need only minor patches.

### 12.2 Value Bridges and Wrapped FC

- **Wrapped FC (wFC)**: During the initial phase, Champions operate a canonical bridge that issues ERC-20 and SPL representations of Foundation Coins on major EVM and Solana networks, similar to early-stage USDC. Each wFC token is 1-for-1 backed by locked FC on the Foundation layer.
- **Atomic swap**: Hashed-time-lock contracts still enable trust-minimized FC – foreign-token exchanges, but wrapped FC offers immediate composability with DeFi protocols that already recognise ERC-20 or SPL assets.
- **Roadmap to native bridges**: Once light-client proofs mature on host chains, the multi-sig bridge will sunset and wFC issuance will migrate to trustless contracts that read Foundation layer headers directly.

### 12.3 Public-Goods Indexes

Services such as **II-Commons** and **II-Thought** expose OpenAPI endpoints for dataset search, model discovery, and licence look-ups. Any chain or dApp can embed these endpoints and earn PoB receipts when they curate or extend the indexes.

### 12.4 Domain Stacks

Specialized stacks apply domain-specific governance while inheriting Foundation layer finality:

- **II-Medical** layers HIPAA-compatible data flags and stricter inference thresholds.
- **II-Legal** embeds citation graphs and licence proofs for every precedent.
- **II-Education** links answers to national learning standards.

### 12.5 Developer Tooling

Reference SDKs (Rust, Python, TypeScript), model-runner containers that auto-emit PoB telemetry, and CLI scaffolds for Common-Ground projects make it easy to build on the stack.

### 12.6 External Compute Escrow (Future Work)

Hooks are reserved for a cross-chain compute escrow: an L2 can borrow idle GPU cycles from Champions, and the resulting PoB receipts flow back to the Foundation layer. The feature is optional in v1 but illustrates how benefit-minting can extend beyond the core network.

With light-client proofs, wrapped FC for immediate DeFi composability, and clear paths for domain and tooling extensions, the Intelligent Internet invites other chains and applications to integrate rather than reinvent.

## Conclusion

The Intelligent Internet is built on a simple equation: verifiable public benefit in, intelligence-backed capital out. By tethering every newly minted Foundation Coin to a cryptographically provable unit of socially useful work, the network aligns scarcity with value, turns open data and models into dependable public infrastructure, and provides a regulatory-grade audit trail for the industries that carry our collective cognitive load.

The technical foundations are laid: a Bitcoin-derived ledger upgraded with one-round BFT finality, Proof-of-Benefit receipts, Anchor-Sets for data provenance, and an agent-centric orchestration model that splits complex tasks into auditable modules. Governance moves in measured steps, from a guarded launch with twelve National Champions and Sentinel oversight to an Oracle Council that tunes parameters on-chain. Wrapped FC gives builders immediate composability on existing chains while light-client bridges mature.