Windfall

A Topocurrency Protocol for Spatially-Routed AI Compute

Patrick Rawson
Ecofrontiers

pat@ecofrontiers.xyz

Version 0.2 — March 19, 2026

Abstract

We introduce **Windfall**, the first *topocurrency* protocol—a system in which the cost, routing, and settlement of economic activity are programmatically modified by verified geographic location. Windfall applies this framework to AI compute, providing an OpenAI-compatible inference gateway that routes requests across a decentralized mesh of nodes to sites where energy is cheapest, cleanest, or most likely to be curtailed. The protocol reads real-time energy oracle data, constructs a continuously-updated locational cost surface, resolves caller identity via onchain registries (ERC-8004, Basenames), and settles payments onchain via Base L2 with verifiable location attestations through Astral Protocol. A semantic cache deduplicates repetitive agent workloads at zero marginal cost, and a graduated free tier rewards verifiable onchain identity. In doing so, Windfall transforms over 100 TWh of annually curtailed clean energy into a productive compute substrate, captures a 57% cost spread between high-congestion and low-congestion zones, and provides AI agents with a native economic primitive for spatially-intelligent infrastructure decisions. This paper formalizes the topocurrency concept, presents Windfall's technical architecture and economic model, and situates the protocol within the broader opportunity to build money that encodes—rather than erases—spatial information.

Contents

1 Introduction

The global data center sector consumed approximately 460 TWh of electricity in 2024 and is projected to reach 945 TWh by 2030—a 15% compound annual growth rate, four times faster than total electricity demand growth [IEA, 2024]. Simultaneously, over 100 TWh of clean energy was curtailed worldwide in the same year: generation that grids could not absorb [LBNL, 2024]. These two facts describe a spatial mismatch of historic proportions. Compute concentrates where infrastructure already exists; clean energy goes uncaptured where it does not.

This mismatch is not a failure of technology. The energy exists. The compute demand exists. What is missing is a *coordination mechanism* that can match geographically flexible demand with locationally fixed supply in real time.

Existing approaches are insufficient. Bilateral Power Purchase Agreements (PPAs) take 6–12 months to negotiate [PPA Industry Report, 2024]. Carbon-aware computing achieves only modest reductions (approximately 5%) because cloud providers cluster in regions with similar grid mixes [Radovanović et al., 2022]. Hyperscaler pricing obscures spatial cost variation through uniform regional pricing. Meanwhile, the neo-cloud market—CoreWeave, Lambda, Crusoe, and bare-metal providers competing on price transparency—is estimated at $24 billion in 2025 and growing at over 40% annually [Neo-cloud Market Report, 2025], creating exactly the price-exposed compute supply that a spatial routing mechanism requires.

Windfall is a protocol that solves this coordination failure. It is an instance of a broader class of systems we call *topocurrencies*: cryptocurrencies whose properties—cost, routing, settlement terms, and governance rights—are programmatically modified by verified geographic location. Where conventional money compresses away spatial information, a topocurrency encodes it. Where conventional markets require bilateral negotiation to match supply and demand across space, a topocurrency automates this matching through a shared, continuously-updated price surface.

This paper proceeds as follows. Section 2 formalizes the topocurrency concept. Section 3 quantifies the spatial compute mismatch. Section 4 presents Windfall's technical architecture—including the energy oracle, spatial router, onchain identity layer, semantic cache, settlement system, and attestation pipeline. Section 5 describes the economic model and token design space. Section 6 addresses security, verification, and trust assumptions. Section 7 positions Windfall within the competitive landscape. Section 9 outlines the development roadmap. Section 12 concludes.

2 Topocurrencies: Money That Knows Where It Is

2.1 Definition

Definition 1 (Topocurrency)**.** *A* ***topocurrency*** *is a cryptocurrency whose properties—value, utility, governance rights, issuance rate, or transaction rules—are programmatically modified by geographic location through verifiable spatial data. Formally, a topocurrency is a tuple* $\mathcal{T} = (\mathcal{M}, \mathcal{L}, \mathcal{O}, \mathcal{G})$ *where:*

- $\mathcal{M}$ *is a monetary mechanism (fee modulation, minting, yield distribution, etc.),*
- $\mathcal{L}$ *is a location proof system producing verifiable spatial claims,*
- $\mathcal{O}$ *is an oracle layer mapping physical-world state to onchain data,*
- $\mathcal{G}$ *is a governance structure determining mechanism parameters.*

The term describes “money that knows where it is”—encoding spatial information as a first-class property of economic transactions rather than an externality to be negotiated around.

2.2 The Problem Topocurrencies Address

Current monetary systems—both fiat and crypto—compress away geographic information, producing three structural failures:

1. **Spatial inequality.** Banks systematically channel deposits from peripheral communities to urban centers where collateral values are higher [De Salvo, 2017]. Cryptocurrencies, by being location-agnostic, reproduce or exceed this inequality pattern.
2. **Ecological blindness.** Ecosystem services are valued at $125–145 trillion annually [Costanza et al., 2014], yet because they are spatially specific, they remain unmeasured in financial systems that erase geography.
3. **Coordination failure.** Spatially differentiated environmental pricing is welfare-optimal [Montgomery, 1972], yet the welfare cost of spatially uniform pricing exceeds $70 billion per year in U.S. air quality regulation alone [Hollingsworth et al., 2022].

Topocurrencies address these failures by making location a programmable parameter of economic activity.

2.3 The Spatial Finance Stack

Three technologies compose into something qualitatively new:

1. **Programmable money** (blockchain): Deterministic execution of conditional logic over value transfer.
2. **Verifiable spatial data** (location proofs): Cryptographically signed claims about where events occur, verified without trusting a single authority.
3. **Autonomous economic actors** (AI agents): Software agents operating at the timescale necessary to read oracle data, evaluate spatial options, and execute transactions.

No single technology is sufficient. Programmable money without spatial data is location-blind. Spatial data without programmable money requires trusted intermediaries. Both without AI agents cannot operate at the temporal resolution (sub-hourly) that real-time spatial arbitrage demands. ERC-8004 (mainnet January 2026) establishes onchain registries for autonomous agents with identity, reputation history, and validation status [ERC-8004, 2026], providing the identity layer that makes topocurrency contracts machine-addressable.

2.4 Mechanism Taxonomy

Topocurrencies decompose into eight mechanisms, each combining a monetary instrument with a spatial primitive:

Mechanism	Description	Spatial Primitive
Fee modulation	Same token, different costs by location	`within()`
Stablecoin capture	External stablecoins locked into geospatial contracts	`contains()`
Service-backed minting	New tokens created when verified outcomes occur at verified locations	`within()`
Spatial yield	Yield allocated by holder location, not just holdings	`within()`
Location-contingent vesting	Tokens unlock only with sustained verified presence	`distance()`
Territorial revenue sharing	Geofenced activity auto-redistributes to local stakeholders	`contains()`
Spatial matching	Two-sided platform matching flexible demand with fixed supply	`within()`
Spatial bonding curves	Token price as a function of geographic demand density	`density()`

Table 1: Topocurrency mechanism taxonomy. Windfall implements spatial matching with fee modulation.

2.5 Spatial Flow Patterns

Value in topocurrency systems flows through geography via six patterns:

- **Global**: No geographic constraint (compute, carbon credits).
- **Place-bound**: Circulates within territory (bioregional currency, tourism).
- **Flow-dependent**: Follows physical paths—water, nutrients, air (watershed services).
- **Commodity**: Travels with physical goods (supply chain provenance).
- **Proximity**: Decays with distance from infrastructure (corridors, networks).
- **Access-gated**: Available only to verified zone participants (humanitarian aid).

Windfall operates in the **global** pattern: AI compute workloads are locationally flexible, and value flows to whichever zone offers the optimal cost surface at any given moment.

3 The Spatial Compute Problem

3.1 The Energy–Compute Mismatch

The AI compute buildout is overwhelming electrical grids. PJM Interconnection's December 2025 capacity auction fell 6,517 MW short of projected demand, with data centers responsible for 94% of the growth projection [PJM, 2025]. Data center electricity consumption is projected to double from 460 TWh (2024) to 945 TWh by 2030 [IEA, 2024].

Simultaneously, clean energy goes wasted at scale:

- CAISO (California) recorded 1,180 hours of negative electricity prices in 2024—13% of all hours [CAISO, 2024].
- ERCOT (Texas) experiences regular negative pricing during wind ramp events lasting approximately 45 minutes.
- Global curtailment exceeded 100 TWh in 2024, representing generation that was available but could not be absorbed by the grid [LBNL, 2024].

The fundamental problem is not energy scarcity but spatial misallocation. Clean energy is abundant in locations without compute infrastructure; compute demand concentrates in locations without clean energy surplus.

3.2 Why Existing Solutions Fail

Bilateral PPAs take 6–12 months to negotiate and are structurally unable to capture transient opportunities—a wind ramp lasting 45 minutes or a negative-price window opening at 2 AM.

Carbon-aware computing achieves only ∼5% carbon reductions because major cloud providers are clustered in regions with similar grid mixes [Radovanović et al., 2022]. Sukprasert et al. tested three years of hourly carbon data across 123 cloud regions and confirmed that geographic load shifting achieves reductions "of the order of 5%" [Sukprasert et al., 2024]. Carbon intensity is too weak a routing signal for actors not already motivated by it—and most compute buyers are motivated by cost, not carbon.

Hyperscaler pricing obscures spatial cost variation. AWS, GCP, and Azure control the majority of cloud compute and use uniform regional pricing that hides the underlying energy cost surface. Spatial cost arbitrage is impossible when the price surface is invisible.

Total locational cost is the comprehensive routing signal:

$$C_{\text{total}}(z,t) = C_{\text{energy}}(z,t) + C_{\text{cool}}(z,t) + C_{\text{water}}(z,t) + C_{\text{cong}}(z,t) \tag{1}$$

where z denotes a geographic zone, t denotes time, and the components are:

- $C_{\text{energy}}(z,t)$: local electricity price (\$/MWh), potentially negative during curtailment,
- $C_{\text{cool}}(z,t)$: energy overhead for thermal management, a function of Power Usage Effectiveness (PUE) and ambient temperature,
- $C_{\text{water}}(z,t)$: water withdrawal and treatment costs,
- $C_{\text{cong}}(z,t)$: transmission congestion from Locational Marginal Pricing (LMP).

3.3 Quantifying the Opportunity

The cost spread between high-congestion and low-congestion zones is substantial:

Zone	C_{energy}	C_{cool}	C_{water}	C_{cong}	C_{total}
Iceland	\$42.00	\$3.00	\$0.50	\$0.00	**\$45.50/MWh**
Finland	\$33.00	\$4.00	\$0.50	\$1.00	**\$38.50/MWh**
N. Virginia	\$70.00	\$25.00	\$3.00	\$8.00	**\$106.00/MWh**

Table 2: Illustrative total locational compute cost by zone. The Iceland–Virginia spread is 57%.

For a 10 MW facility operating 8,000 hours/year (80 GWh), routing from Northern Virginia to Iceland yields approximately **\$4.8 million in annual savings**. At the neo-cloud market scale (\$24B, growing 40%+ annually), even modest spatial optimization over a fraction of workloads represents a multi-billion dollar opportunity.

3.4 Flexibility Constraints

Honest assessment of what can move:

- The commonly cited 30–50% "flexible AI workload" estimate overstates what is practically movable.
- Frontier models with petabytes of co-located weights and thousands of GPUs cannot pause and relocate quickly.
- Data sovereignty constraints (GDPR, national security) restrict where workloads can execute.
- The locationally flexible fraction of total data center energy is realistically **10–20% today**, growing as AI inference (vs. training) becomes the dominant workload.
- Inference is inherently more flexible than training: stateless, latency-tolerant for non-interactive use cases, and scales horizontally across sites.

3.5 Why Onchain Settlement

A skeptic observes that Google already routes compute spatially without any cryptocurrency. The argument for onchain settlement rests on four specific market failures that a centralized platform cannot efficiently solve:

1. **Transaction costs.** Bilateral PPAs take 6–12 months to negotiate and require specialized legal counsel [PPA Industry Report, 2024]. For real-time spatial matching at 5-minute oracle resolution, bilateral contracting is structurally impossible. Smart contracts have demonstrated order-of-magnitude reductions in settlement friction: JP Morgan's Quorum platform cut interbank settlement costs by 50–70%, and a controlled study of supply chain contracting found a 76% reduction in per-transaction fees with settlement time compressed by 70% [Smart Contract Efficiency, 2025].
2. **Coordination failure.** Matching flexible compute demand with spatially distributed renewable supply is a multilateral matching problem [Roth, 2002]. Bilateral contracting produces "application waste" and "screening waste": each negotiation ignores its externalities on all other potential matches. A *centralized* clearinghouse solves coordination but introduces rent extraction—intermediaries capture a larger share of gains from trade. A decentralized clearinghouse (smart contract + oracle) avoids this: the matching rules are public, execution is automatic, and no single intermediary can extract rent by controlling the matching rules.
3. **Hold-up.** Co-location investments are relationship-specific—once a data center builds next to a wind farm, it is vulnerable to renegotiation. This is the hold-up problem Williamson formalized: asset specificity combined with incomplete contracts enables opportunism [Williamson, 2009]. Smart contracts function as credible commitment devices—immutable terms, automatic execution—resolving hold-up without courts or renegotiation.
4. **Composability.** A centralized platform can route compute. What it cannot do is compose routing with adjacent financial functions in a single transaction. A facility operator needs not just a matched workload, but insurance against curtailment intermittency, hedging on energy price movements, and capital to build the facility. In traditional markets, each requires a separate counterparty, separate collateral, and separate settlement. Onchain, these compose: the same USDC settlement that pays for routing can collateralize a curtailment insurance position; onchain revenue history becomes transparent proof-of-income for lending protocols financing new renewable-adjacent facilities.

The honest counterpoint. Centralization is better for high-frequency, low-latency operations where millisecond execution matters. Abadi and Brunnermeier formalize the blockchain trilemma: a consensus algorithm cannot simultaneously achieve fault-tolerance, resource-efficiency, and full transferability [Abadi and Brunnermeier, 2018]. The argument is not that oracle aggregation, route optimization, or workload scheduling must be onchain—these are computation-heavy tasks better handled offchain. It is that the *settlement layer* and *commitment mechanism* benefit from decentralization—public rules, automatic execution, and contestable intermediary functions—while the compute routing itself uses offchain optimization that settles onchain.

4 Windfall Protocol Architecture

4.1 Overview

Windfall is a spatially-routed inference gateway deployed on Base L2. It implements two topocurrency mechanisms—**spatial matching** and **fee modulation**—to route AI inference requests to compute nodes based on a real-time locational cost surface. The gateway is OpenAI-compatible: any agent or application that speaks the `/v1/chat/completions` interface can point at Windfall as a drop-in replacement and gain spatial routing automatically.

The architecture comprises eight components:

1. **Energy Oracle**: Polls real-time grid data (energy price, carbon intensity, renewable percentage, curtailment status) across all nodes via the Electricity Maps API every 5 minutes.
2. **Cost Surface**: Aggregates oracle data into a continuously-updated scalar field $C_{\text{total}}(z, t)$ across the node mesh.
3. **Spatial Router**: Selects the optimal node for each inference request based on the caller's routing preference. Default mode is `greenest`.
4. **Inference Proxy**: Forwards requests to the selected node, proxying to LLM providers (OpenRouter) with location-aware metadata.
5. **Identity Layer**: Resolves the caller's onchain identity—anonymous wallet, Basename holder, or ERC-8004 registered agent—and assigns a trust tier that determines free-tier allocation, rate limits, and access privileges.
6. **Semantic Cache**: Deduplicates identical or near-identical requests via SHA-256 content hashing with a 1-hour TTL, scoped per caller. Cache hits return instantly at zero cost, reducing both latency and expenditure for repetitive agent workflows.
7. **Settlement Layer**: Handles payment across four paths—API key balance (topped up via Stripe card payment, USDC deposit, or ETH deposit), direct onchain transfer (USDC or ETH on Base), x402 machine-payable HTTP protocol, and a graduated free tier—and publishes verifiable attestations via the Ethereum Attestation Service (EAS) on Base.
8. **Attestation Layer**: Every routed inference produces a verifiable location-energy attestation anchored onchain via EAS, using the Astral Protocol location proof schema.

4.2 Energy Oracle

The energy oracle maintains a real-time view of grid conditions across all nodes in the Windfall mesh. The current implementation polls the Electricity Maps API (v3) every 5 minutes for each node's grid zone, retrieving:

- Power consumption breakdown by source (wind, solar, hydro, biomass, geothermal, nuclear, fossil),
- Real-time carbon intensity (gCO_2eq/kWh),
- Derived renewable percentage and estimated spot price.

A curtailment heuristic flags zones where renewable percentage exceeds 85% and carbon intensity falls below 50 gCO_2/kWh—conditions that empirically correspond to grid oversupply and potential negative pricing.

The oracle produces an `EnergyCostSurface` object: a map from node identifiers to energy data, annotated with the cheapest and greenest nodes at each polling interval. Staleness detection triggers fallback to static baseline prices if oracle data exceeds a 15-minute threshold.

Definition 2 (Cost Surface)**.** *The cost surface is a time-varying scalar field over the set of nodes* $\mathcal{N}$*:*

$$S : \mathcal{N} \times \mathbb{R}^+ \to \mathbb{R}^4$$

mapping each node $n \in \mathcal{N}$ *at time* t *to a vector* $(C_{energy}, C_{cool}, C_{water}, C_{cong})$*. The routing function selects* $n^* = \arg\min_{n \in \mathcal{N}_{healthy}} f(S(n,t))$ *where* f *is a preference-weighted aggregation and* $\mathcal{N}_{healthy}$ *is the set of nodes passing health checks.*

Roadmap: Oracle expansion. The production oracle will integrate direct spot price feeds (Nord Pool for Nordic zones, ERCOT nodal prices for Texas, CAISO LMP data for California), weather APIs for cooling load estimation, and water stress indices. Oracle resolution will increase from 5-minute polling to sub-minute websocket streams where ISOs provide them.

4.3 Spatial Router

The spatial router implements three routing modes:

`cheapest` Minimizes energy cost: $n^* = \arg\min_n C_{\text{energy}}(n, t)$.

`greenest` Minimizes carbon intensity: $n^* = \arg\min_n \mathrm{CI}(n, t)$.

`balanced` Minimizes a Pareto-weighted score. Both price and carbon are normalized to $[0, 1]$ across healthy candidates:

$$\text{score}(n) = w_p \cdot \frac{C_{\text{energy}}(n,t) - C_{\text{energy}}^{\min}}{C_{\text{energy}}^{\max} - C_{\text{energy}}^{\min}} + w_c \cdot \frac{\mathrm{CI}(n,t) - \mathrm{CI}^{\min}}{\mathrm{CI}^{\max} - \mathrm{CI}^{\min}} \tag{2}$$

where $w_p = w_c = 0.5$ by default. The caller can override these weights, enabling AI agents to express nuanced preferences (e.g., $w_p = 0.8, w_c = 0.2$ for cost-sensitive batch jobs; $w_p = 0.1, w_c = 0.9$ for carbon-committed organizations).

The **default routing mode is `greenest`**—minimizing carbon intensity—reflecting the protocol's design conviction that ecological optimization should be the path of least resistance. Callers who want cost optimization must explicitly request it, inverting the usual default where cost comes first and sustainability is opt-in.

Node health is verified via periodic liveness probes with latency tracking. If all peer nodes fail health checks, the router falls back to the local node, ensuring availability. Inter-node proxying uses authenticated headers (`X-Proxied-From`, `X-Payment-Verified`) so that the entry node handles payment and the executing node performs inference without double-charging.

4.4 Settlement

Windfall settles on **Base** (Coinbase L2, Ethereum-secured). Four payment paths are supported, designed to serve the full spectrum from first-time explorers to high-volume autonomous agents:

1. **API key balance**: The primary payment method. Callers generate an API key tied to a wallet address and top up its balance via three channels:
    - **Stripe card payment** ($5, $10, $25, or $50 increments) for onboarding users without crypto,
    - **USDC deposit** on Base, auto-credited within 60 seconds via a deposit watcher polling every 30 seconds,
    - **ETH deposit** on Base, converted at the CoinGecko spot rate and auto-credited via the same watcher cadence.

    Each inference request deducts from the key balance. Usage, savings, and routing history are queryable via the API and a browser-based dashboard.

2. **Direct onchain transfer**: A caller passes a transaction hash (`X-Payment-TX` header) for a USDC or ETH transfer on Base. The gateway verifies the transfer onchain before executing inference. Replay protection prevents reuse of transaction hashes.

3. **x402**: The machine-payable HTTP protocol (Coinbase, 2026) enabling AI agents to pay for API calls natively via HTTP headers without pre-established accounts.

4. **Graduated free tier**: Rather than a flat allocation, free requests scale with onchain identity:
   - **Anonymous**: 25 free requests (no wallet connected),
   - **Wallet**: 50 free requests (wallet address provided),
   - **Basename**: 100 free requests (Coinbase Basename resolved onchain),
   - **ERC-8004 agent**: 100 free requests (registered in the Base agent registry).

   The graduated structure rewards callers who establish verifiable onchain identity—the same signal the protocol uses for trust and reputation.

Usage tracking persists in a SQLite database with WAL journaling. Deposit watchers for both USDC and ETH poll at 30-second intervals, verifying transfer amounts onchain before crediting key balances.

Fee structure. Base inference: \$0.004/request. Premium models (Claude, GPT-4, o1): \$0.008/request. A 10% green surcharge applies when the caller explicitly selects `greenest` mode, funding a protocol sustainability reserve. **Cache hits are free**: when the semantic cache returns a previously computed response, no inference is executed and no fee is charged—reducing both cost and energy consumption for repetitive agent workflows.

4.5 Onchain Identity

Windfall resolves every caller's onchain identity at request time, producing a trust tier that governs free-tier allocation and future access privileges. The resolution checks three sources in sequence:

1. **ERC-8004 registry**: If the wallet address is registered in the Base agent registry (mainnet January 2026), the caller is classified as a verified autonomous agent.
2. **Basename**: If the wallet resolves to a Coinbase Basename via onchain lookup, the caller is classified as an identified human or organization.
3. **Wallet**: If neither registry returns a match, the caller is classified by wallet address alone.

The identity tier is stored alongside the API key and informs the graduated free tier. The design principle is that *verifiable onchain identity should be rewarded*, not required: anonymous callers can still use the protocol, but identified callers receive more generous defaults. As the protocol matures, identity tiers will gate governance participation and spatial reward eligibility.

4.6 Semantic Cache

Windfall maintains a semantic response cache that deduplicates identical or near-identical inference requests. Cache keys are SHA-256 hashes of the normalized message content, model identifier, and caller scope (API key, wallet address, or anonymous session). Cached responses have a 1-hour TTL and are automatically expired every 10 minutes.

Cache hits return instantly at zero cost—no inference is executed, no energy is consumed, and no fee is charged. For AI agent workflows that repeatedly query similar prompts (status checks, environmental data lookups, routing queries), this can eliminate a substantial fraction of inference volume. The cache is scoped per caller to prevent information leakage between users.

The economic implication is that the protocol's marginal cost curve is nonlinear: the first request for a given query costs \$0.004–\$0.008, but subsequent identical requests cost nothing. This incentivizes agents to route through Windfall for repetitive workloads, building habit and volume even before the spatial routing benefit is the primary draw.

4.7 Attestations

Every inference routed through Windfall produces a verifiable attestation via the **Ethereum Attestation Service** (EAS) on Base (contract `0x4200...0021`), using the Astral Protocol location proof schema (UID: `0xba41...24e2`). Location data is encoded as GeoJSON Points (EPSG:4326). Attestations are non-revocable and addressed to the zero address (public).

Attestations are **batched** for gas efficiency: individual inference events queue locally and flush every 5 minutes or when 50 events accumulate, whichever comes first. Each batched attestation encodes:

- Timestamp of inference execution,
- Node identifier and geographic coordinates (latitude, longitude) as GeoJSON Point,
- Energy price, carbon intensity, renewable percentage, and curtailment status at time of execution,
- Model used and SHA-256 hash of the response,
- Batch request count (for aggregated attestations),
- Gateway version and memo metadata in JSON format.

These attestations serve three functions:

1. **Verifiability**: Any third party can confirm where, when, and under what energy conditions a given inference was executed. Each attestation is browsable on EASScan.
2. **Composability**: Downstream protocols can gate actions on Windfall attestation data—e.g., a carbon credit protocol verifying that an AI agent's compute was renewable-powered.
3. **Reputation**: Over time, attestations form an onchain history of an agent's energy preferences and routing decisions, enabling reputation-based access and pricing.

Astral Protocol integration. Windfall's attestation layer is designed to integrate with Astral Protocol's location proof framework. Astral provides a standardized schema for location claims, location stamps (independent evidence sources across seven categories: authority, social, near-field machine, network machine, sensor data, delegated, and legal), and location proofs (multi-stamp bundles producing a multidimensional **CredibilityVector**—evaluating location evidence across independent dimensions of spatial accuracy, temporal alignment, stamp validity, and source independence). An application-specific weighting function collapses the vector into a single **CredibilityAssessment** score from 0–1 that contracts can threshold against; the vector itself is the richer artifact [Astral, 2025]. The integration path:

- **Phase 1 (current)**: EAS attestations with self-reported node coordinates.
- **Phase 2**: Astral Location Protocol records wrapping EAS attestations, adding cryptographic signatures over spatial data.
- **Phase 3**: Multi-stamp location proofs combining network-machine evidence (IP geolocation, cell triangulation), sensor data (power meter readings), and delegated verification (facility operator co-signatures) into location proofs with quantified credibility scores.

Zero-knowledge point-in-polygon proofs (ZKMaps) will enable Windfall to verify that inference executed within a specific grid zone without revealing exact facility coordinates—critical for operators who consider precise location commercially sensitive [ZKMaps, 2025].

4.8 System Diagram

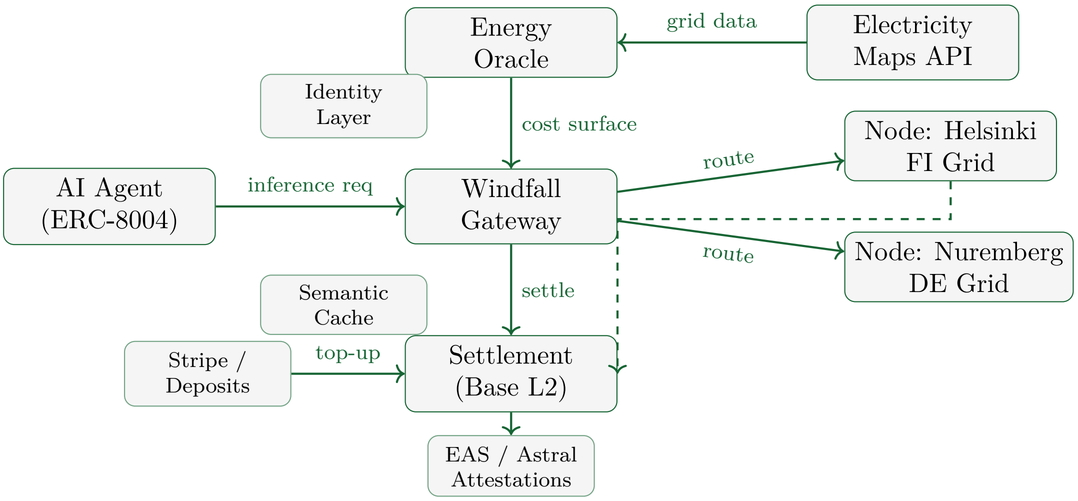

Figure 1: Windfall protocol architecture. AI agents submit inference requests with routing preferences. The gateway resolves caller identity, checks the semantic cache, consults the energy oracle's cost surface, routes to the optimal node, settles payment on Base (via API key balance, onchain transfer, x402, or free tier), and publishes location-energy attestations via EAS.

5 Economic Model

5.1 Market Structure

Windfall is a **two-sided matching platform**, not merely a pricing mechanism:

- **Demand side**: AI agents seeking inference compute, varying in latency tolerance, cost sensitivity, carbon commitment, and data sovereignty requirements.
- **Supply side**: Compute node operators at diverse geographic locations, varying in energy cost, renewable mix, cooling efficiency, and available capacity.

Cross-side network effects drive adoption: more flexible workloads attract node operators to join the mesh in renewable-rich locations; more node diversity attracts compute buyers who benefit from better spatial options. This is the structure Rochet and Tirole describe in their foundational work on two-sided platforms: an intermediary that creates value by connecting two distinct groups whose participation benefits the other side [Rochet and Tirole, 2003]. The spatial dimension creates natural geographic diversification incentives rather than winner-take-all concentration—each new node in a distinct grid zone adds routing options that no existing node can substitute.

5.2 Fee Model

Revenue flows through four channels:

1. **Per-request fees**: \$0.004 (standard) to \$0.008 (premium) per inference request. Competitive with direct OpenRouter pricing plus the spatial optimization benefit. Cache hits incur no fee—the protocol earns nothing on deduplicated requests, trading short-term revenue for agent retention.
2. **Green surcharge**: 10% premium on `greenest`-mode requests, directed to a sustainability reserve that funds oracle infrastructure and node onboarding in renewable-rich zones.

3. **Spread capture**: When energy price differentials are large (curtailment events, negative pricing windows), Windfall captures a fraction of the cost savings as protocol revenue, sharing the remainder with the compute buyer.

4. **Fiat on-ramp**: Stripe card payments provide a bridge for users without crypto, with standard payment processing applied to top-up amounts. This lowers the barrier to entry and funds the same API key balance that crypto deposits credit.

5.3 Saturation Dynamics

A critical property of the cost surface is **self-correcting saturation**: as compute concentrates in a low-cost zone, energy prices rise, congestion increases, cooling infrastructure saturates, and the zone's cost advantage diminishes. The oracle reflects this in real time, and routing shifts to the next-cheapest zone.

Proposition 1 (Saturation equilibrium)**.** *Under the assumption that oracle data is timely and accurate, the Windfall routing mechanism converges to a spatial equilibrium in which no compute buyer can reduce cost by unilateral relocation. Formally, at equilibrium:*

$$C_{total}(n_i, t) = C_{total}(n_j, t) + \epsilon_{ij} \quad \forall\, n_i, n_j \in \mathcal{N}_{active} \tag{3}$$

where ϵ_{ij} captures non-energy switching costs (latency, data transfer, sovereignty constraints) between nodes i and j.

This means Windfall does not simply redirect all compute to one cheap location; it distributes load across the mesh in proportion to each zone's sustainable cost advantage.

5.4 Token Design Space

Windfall currently settles in USDC and ETH. A protocol-native token is under consideration and would serve several potential functions:

1. **Node operator staking**: Operators stake tokens to join the mesh, creating skin-in-the-game for honest location reporting and uptime. Slashing conditions penalize false attestations.

2. **Governance**: Token holders vote on mechanism parameters: routing weights, fee schedules, node onboarding criteria, oracle source selection, and green surcharge allocation.

3. **Fee discounts**: Token holders receive discounted inference rates, aligning long-term protocol usage with token demand.

4. **Spatial rewards**: Bonus token emissions for operators in zones with high renewable percentage or active curtailment, incentivizing node deployment where it produces the greatest environmental benefit.

5. **Insurance underwriting**: Token holders underwrite insurance against curtailment risk—guaranteeing compute buyers that if a node goes offline due to grid instability, they can claim against the insurance pool and be re-routed seamlessly.

The design principle is that any native token must create utility that USDC/ETH alone cannot provide. We explicitly reject token models that exist solely to create speculative demand. The token design will be finalized through community governance and published in a subsequent tokenomics paper.

6 Security, Verification, and Trust Assumptions

6.1 Threat Model

Windfall's security model addresses four primary threats:

1. **Location spoofing**: A node operator claims to be in a low-cost zone while actually operating in a high-cost zone, capturing routing preference without providing the cost benefit.
2. **Oracle manipulation**: Corrupted or stale oracle data causes suboptimal routing, either wasting compute buyer funds or enabling arbitrage by informed insiders.
3. **Sybil attacks**: An operator registers multiple virtual "nodes" at advantageous locations to capture disproportionate routing volume.
4. **Data integrity**: If token value or routing preference depends on environmental metrics, operators have incentives to fabricate favorable data.

6.2 Mitigation Strategies

Location verification follows the Astral Protocol model of multi-signal composition [Astral, 2025]. No single location signal is trusted. Instead, independent evidence sources are combined:

- Network-level: IP geolocation, BGP routing analysis, latency triangulation from multiple peers.
- Physical: Power meter data co-signed by the facility operator and the grid operator.
- Sensor: Ambient temperature readings cross-referenced with weather APIs.
- Legal: Facility lease agreements or utility bills with geographic addresses.

The key insight from location proof theory is that the cost of forging multiple independent signals simultaneously exceeds the economic benefit of spoofing for most transaction values [Astral, 2025]. Windfall's credibility threshold scales with transaction value: low-value inference requests accept lower-certainty location proofs; high-value routing commitments require stronger multi-stamp verification.

Oracle reliability uses defense-in-depth:

- Multi-source aggregation: Cross-referencing Electricity Maps, direct ISO feeds, and weather data.
- Staleness detection: 15-minute threshold triggers fallback to conservative baseline prices.
- Economic staking: Oracle providers stake tokens that are slashed for provably incorrect data.
- Circuit breakers: Extreme price movements ($>3\sigma$ from rolling average) trigger manual review before routing decisions execute.

Sybil resistance combines staking requirements (capital cost per node) with physical verification (each node must demonstrate independent power metering and distinct network paths).

Data integrity requires adversarial separation between data production and economic benefit. Node operators whose compensation depends on environmental metrics must not control the sensors producing those metrics.

Operational hardening. The deployed gateway implements defense-in-depth at the application layer:

- **Rate limiting**: Inference (60/min), key creation (5/hour), authentication (5 per 15 min), general API (120/min), and contact forms (3/hour).
- **CORS**: Restricted to known origins (`windfall.ecofrontiers.xyz`, `ecofrontiers.xyz`).
- **Transaction replay protection**: Used transaction hashes are tracked in-memory and rejected on resubmission.

- **Key security**: API keys are SHA-256 hashed at rest; plaintext keys are returned only at creation time and never logged.
- **Session management**: Admin sessions expire after 24 hours; wallet authentication uses Sign-In With Ethereum (SIWE) with nonce-based replay protection.
- **Data retention**: Contact submissions purged after 12 months, request logs anonymized after 30 days, inactive API keys deleted after 12 months.

6.3 Trust Assumptions

We are explicit about what Windfall trusts:

- **Base L2 consensus**: Settlement assumes Base is live and honest. Base inherits Ethereum's security through optimistic rollup fraud proofs.
- **Oracle data feeds**: Grid pricing data is assumed accurate within the polling interval. This is the weakest link—oracle corruption could misdirect routing.
- **LLM providers**: Inference quality depends on upstream providers (OpenRouter → model providers). Windfall verifies response hashes but cannot verify inference quality.
- **Node operators**: Operators are assumed to execute inference honestly but are not trusted on location claims without multi-signal verification.

7 Competitive Landscape

7.1 Positioning: Infrastructure, Not Cloud

Windfall is **complementary** to existing compute providers, not competitive with them. The protocol routes requests *to* neo-cloud and bare-metal providers; it does not operate data centers or serve inference directly.

Category	**Examples**	**Windfall relationship**
Hyperscalers	AWS, GCP, Azure	Opaque pricing hides spatial cost variation. Windfall operates in the price-transparent segment they do not serve.
Neo-cloud	CoreWeave, Lambda, Crusoe	Supply-side participants. Their price transparency creates the cost surface Windfall routes over.
Decentralized compute	Akash, Render, io.net	Adjacent but architecturally distinct. These are general-purpose compute marketplaces; Windfall is a spatial routing layer that could sit above them.
Carbon-aware computing	Google CarbonAware, WattTime	Single-signal (carbon only) routing. Windfall uses multi-signal cost surface including price, cooling, water, and congestion.

Table 3: Competitive landscape. Windfall is infrastructure that routes to existing providers, not a competing provider.

7.2 Moat

Windfall's defensibility comes from three sources:

1. **Network effects**: Each new node improves routing options for all compute buyers; each new buyer increases revenue for all operators. This is the standard two-sided platform moat, but with a spatial dimension that creates natural diversification incentives.
2. **Attestation history**: Over time, the onchain record of spatially-verified inference creates a composable data asset. Protocols building on attestation data (carbon accounting, ESG reporting, AI safety auditing) create lock-in for the attestation standard.
3. **Oracle infrastructure**: The multi-source energy oracle, once built and validated, is expensive to replicate. Integrating direct ISO feeds, weather APIs, and facility-level metering requires domain expertise and data partnerships that compound over time.

8 The Bonded Surface: Broader Vision

Windfall's energy cost surface is a special case of a more general object: a **bonded surface** that assigns ecological price to every coordinate on Earth:

$$P(\text{lat}, \text{lon}, t) = \sum_i w_i \cdot C_i(\text{lat}, \text{lon}, t) \tag{4}$$

where C_i are data layers (grid carbon intensity, water stress, biodiversity indices, soil health, land cover change, air quality, ocean health) and w_i are governance-set weights.

This surface creates a dynamic equilibrium: as economic activity concentrates in a zone, its ecological surplus is consumed—local price rises. Activity redistributes to the next-cheapest zone. No central authority directs flow; the price gradient does the work. The surface self-adjusts.

This realizes Montgomery's (1972) proof that spatially differentiated environmental pricing is welfare-optimal [Montgomery, 1972]—a result policymakers have never operationalized at scale. The welfare cost of spatially uniform pricing is estimated at $70 billion per year for U.S. air quality regulation alone [Hollingsworth et al., 2022], and $310–940 million per year from the SO_2 trading program's failure to incorporate spatial damage ratios.

Windfall's energy cost surface is the first instantiation of this bonded surface for a single dimension (energy). The topocurrency framework envisions its extension to multiple ecological dimensions, governed polycentrically—bioregional compacts setting local weights, cross-jurisdictional agreements for shared resources, and global weights only for true commons.

8.1 Beyond Compute: The Topocurrency Use Case Matrix

Use case	**Mechanism**	**Flow**	**Primitive**	**Status**
Spatial compute	Matching + fees	Global	`within()`	**Prototype**
Carbon pricing	Fee modulation	Global	`within()`	Theoretical
Bioregional ecology	Minting + revenue	Place-bound	`contains()`	Piloted
Watershed services	Spatial yield	Flow-dep.	`within()`	Theoretical
Tourism	Stablecoin capture	Place-bound	`contains()`	Theoretical
Supply chain	Location vesting	Commodity	Origin attest.	Theoretical
Humanitarian	Stablecoin capture	Access-gated	Location proof	Deployed*

Table 4: Topocurrency use case matrix. Windfall is the first purpose-built topocurrency protocol. *WFP Building Blocks moves $555M monthly across 65 organizations using stablecoin capture at scale.

9 Roadmap

Phase	Milestones
Phase 0: Prototype (Q1 2026)	Working gateway with 2 nodes (Helsinki, Nuremberg). Electricity Maps oracle with 5-minute polling and curtailment detection. EAS attestations on Base using Astral schema with batched publishing. OpenAI-compatible inference proxy via OpenRouter. Four payment paths: API key balance (Stripe, USDC deposit, ETH deposit), direct onchain transfer, x402, and graduated free tier (25–100 requests by identity tier). Semantic response cache. Onchain identity resolution (wallet, Basename, ERC-8004). Browser dashboard and admin command center. SIWE authentication.
Phase 1: Expand (Q2–Q3 2026)	5–10 nodes across Nordic, Iberian, and ERCOT zones. Direct ISO price feeds (Nord Pool, ERCOT). x402 machine-payable integration. Astral Location Protocol v1 attestations. Public API and SDK for AI agent frameworks.
Phase 2: Harden (Q4 2026)	Multi-stamp location proofs (Astral Phase 3). Oracle staking and slashing. ZK location verification for operator privacy. Insurance pool for curtailment risk. Token design finalization and governance launch.
Phase 3: Scale (2027)	50+ nodes globally, including CAISO, PJM, Australian NEM, and South American grids. Composable attestation marketplace. Integration with decentralized compute networks (Akash, io.net). Full bonded surface with multi-dimensional ecological pricing.

Table 5: Development roadmap.

10 Governance

Windfall's governance must answer fundamentally political questions: Who sets oracle weights? Who defines zone boundaries? Who determines fee schedules and node onboarding criteria?

We adopt a **polycentric governance** model, drawing on Ostrom's framework for governing common-pool resources [Ostrom, 1990]:

- **Local layer**: Node operators govern operational parameters for their facilities (capacity limits, maintenance windows, local pricing floors).
- **Regional layer**: Operators within a grid zone collectively set zone-level parameters (minimum renewable thresholds, curtailment response policies).
- **Protocol layer**: All stakeholders (operators, compute buyers, token holders) govern global parameters (fee schedules, oracle source selection, routing algorithm weights, new zone onboarding).

Governance rights are initially contribution-weighted: operators who have run nodes longest and processed the most verified inference carry greater weight. As the token model matures, governance transitions to a hybrid of stake-weighted and contribution-weighted voting, preventing pure plutocratic capture.

11 Honest Limitations

1. **Windfall does not disrupt hyperscalers.** AWS, GCP, and Azure control the majority of cloud compute and hide spatial cost variation through uniform regional pricing. They have no structural incentive to expose a real-time cost surface. Windfall operates in the neo-cloud and bare-metal segment—CoreWeave, Lambda, Crusoe, and bare-metal operators competing on price transparency—where the cost surface is visible. If hyperscaler adoption comes, it arrives from the demand side: enterprise buyers who need spatial transparency for carbon reporting under the EU's Corporate Sustainability Reporting Directive (CSRD), which requires Scope 2 emissions disclosure using both location-based and market-based methods.

2. **Curtailment is a finite resource.** Global 2024 curtailment was approximately 100 TWh; locationally flexible data center energy is approximately 40–80 TWh annually. Battery storage deployment is reducing curtailment—CAISO alone added 3.6 GW of battery capacity in 2024—and grid expansion eventually reaches remote generation sites. A compute facility built on the assumption of permanently cheap curtailed energy may find its cost advantage eroding as infrastructure catches up. The oracle-driven pricing mitigates this dynamically: as curtailment falls and prices normalize in a zone, routing shifts elsewhere. But physical infrastructure cannot move.

3. **Cost savings first, emissions second.** The mechanism optimizes on total cost, not carbon. Emissions fall only where cheap zones also happen to be clean zones. That overlap is substantial today—Iceland, Nordic hydro, ERCOT wind—but it is not guaranteed. Knittel et al. (2025) demonstrate that flexible data centers reduce emissions only when regional renewable penetration exceeds approximately 50%; below that threshold, load flexibility can support fossil baseload by filling off-peak valleys, *increasing* emissions [Knittel et al., 2025]. A zone that is cheap because of lax environmental regulation, not renewable abundance, would attract compute under this mechanism just as effectively.

4. **Cost surface dynamics remain partially unmapped.** Our formalization treats C_{total} as exogenous, but spatial general equilibrium effects (energy market feedback, labor market effects, capital reallocation) could amplify or attenuate cost signals. If compute floods Iceland, land prices rise, labor scarcity appears, and the advantage erodes *because* activity arrived. Conversely, agglomeration economies could make popular destinations *cheaper* over time. Which effect dominates is an empirical question requiring spatial general equilibrium modeling—the framework Balboni and Shapiro (2025) identify as the frontier for this kind of analysis [Balboni and Shapiro, 2025].

5. **The oracle is the weakest link.** Data layer integrity depends on underlying feeds. Multi-source aggregation, staking, and circuit breakers mitigate but do not eliminate this risk.

6. **Location proofs are probabilistic, not absolute.** Multi-signal composition raises forgery cost but does not eliminate it. The system's security guarantee is economic (forgery cost exceeds transaction value), not cryptographic.

7. **Governance capture is an unsolved problem.** Who draws zone boundaries and sets oracle weights are political questions with economic consequences. Polycentric governance is our best tool, but no precedent exists for this at the scale topocurrencies envision.

8. **Fee-gaming.** Agents may register in low-fee zones while operating elsewhere. Goodhart's Law applies: once agents understand the cost function, they optimize against it. Continuous mechanism evolution is required. The deeper problem is that any published cost

function becomes a target for adversarial optimization—the same dynamic that plagues SEO, credit scoring, and emissions testing.

12 Conclusion

The spatial mismatch between AI compute demand and clean energy supply is one of the defining infrastructure challenges of the coming decade. Over 100 TWh of clean energy goes uncaptured annually while data centers overwhelm grids in concentrated zones. The cost spread between high-congestion and low-congestion regions exceeds 57%.

Windfall addresses this mismatch through a topocurrency protocol: money that knows where it is, routing AI inference to where energy is cheapest, cleanest, and most available. The protocol combines real-time energy oracles, spatial routing algorithms, onchain settlement on Base L2, and verifiable location attestations through Astral Protocol into a coherent system that replaces 6–12 month bilateral negotiations with real-time, programmatic spatial matching.

More broadly, Windfall is the first instantiation of the topocurrency thesis—the claim that geographic information should be encoded into economic transactions rather than compressed away. The energy cost surface it constructs is a special case of the bonded surface: a continuously-updated scalar field assigning ecological price to geography, enabling markets to internalize spatial externalities that have been systematically ignored.

We build Windfall not because spatial compute is the only application of topocurrencies, but because it is the most immediately addressable—the data exists, the cost differentials are large, the flexible workloads are growing, and the infrastructure (Base, Astral, Electricity Maps, x402) is deployed. It is the proving ground for a broader vision: money that makes geography legible to markets, and markets responsive to geography.

References

Astral Protocol. Notes on location proofs. Technical report, Astral, September 2025.

California Independent System Operator. Annual report on market issues and performance, 2024.

R. Costanza, R. de Groot, P. Sutton, S. van der Ploeg, S. Anderson, I. Kubiszewski, S. Farber, and R. K. Turner. Changes in the global value of ecosystem services. *Global Environmental Change*, 26:152–158, 2014.

J. S. De Salvo. Banking geography and spatial inequality. *Journal of Economic Geography*, 17(4):775–800, 2017.

ERC-8004: Agent registry standard. Ethereum Improvement Proposals, mainnet January 2026.

A. Hollingsworth, I. Rudik, and C. Vidovic. The welfare costs of spatially uniform environmental regulation. Working paper, 2022.

International Energy Agency. Electricity 2024: Analysis and forecast to 2026. IEA, Paris, 2024.

Lawrence Berkeley National Laboratory. Wind and solar curtailment in the United States and globally, 2024.

W. D. Montgomery. Markets in licenses and efficient pollution control programs. *Journal of Economic Theory*, 5(3):395–418, 1972.

Industry analysis of the neo-cloud compute market. Market sizing estimates, 2025.

E. Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, 1990.

PJM Interconnection. 2025/2026 reliability pricing model base residual auction results. December 2025.

Corporate PPA negotiation timelines. Industry report, 2024.

A. Radovanović et al. Carbon-aware computing for datacenters. *IEEE Transactions on Power Systems*, 2022.

J.-C. Rochet and J. Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029, 2003.

A. E. Roth. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70(4):1341–1378, 2002.

The impact of blockchain-enabled smart contracts on firms' operational efficiency. *Journal of Operations Management*, 2025.

T. Sukprasert, A. Souza, N. Bashir, D. Irwin, and P. Shenoy. On the limitations of carbon-aware temporal and spatial workload shifting in the cloud. In *Proceedings of EuroSys '24*, 2024.

O. E. Williamson. Transaction cost economics: The natural progression. Nobel Prize Lecture, December 2009.

J. Abadi and M. Brunnermeier. Blockchain economics. NBER Working Paper 25407, 2018 (revised 2022).

C. R. Knittel, J. R. L. Senga, and S. Wang. Flexible data centers and the grid: Lower costs, higher emissions? NBER Working Paper 34065, 2025.

C. A. Balboni and J. S. Shapiro. Spatial environmental economics. NBER Working Paper 33377, 2025.

ZKMaps. Private onchain location data via zero-knowledge proofs. 2025.

A Protocol Parameters

Parameter	Description	Default
`PRICE_PER_REQUEST`	Base inference fee	\$0.004
`PREMIUM_PRICE_PER_REQUEST`	Premium model fee	\$0.008
`GREEN_SURCHARGE`	Surcharge for `greenest` mode	10%
`ORACLE_POLL_INTERVAL`	Energy data polling frequency	5 min
`ORACLE_STALE_THRESHOLD`	Max oracle data age before fallback	15 min
`FREE_TIER_ANONYMOUS`	Free requests (no wallet)	25
`FREE_TIER_WALLET`	Free requests (wallet connected)	50
`FREE_TIER_BASENAME`	Free requests (Basename holder)	100
`FREE_TIER_ERC8004`	Free requests (registered agent)	100
`ATTESTATION_BATCH_SIZE`	EAS attestation batch size	50
`ATTESTATION_FLUSH_INTERVAL`	Max time before attestation publish	5 min
`CACHE_TTL`	Semantic cache time-to-live	1 hour
`DEPOSIT_POLL_INTERVAL`	USDC/ETH deposit watcher cadence	30 sec
`DEFAULT_ROUTING_MODE`	Routing mode when unspecified	`greenest`
`BALANCED_WEIGHT_PRICE`	Price weight in balanced mode	0.5
`BALANCED_WEIGHT_CARBON`	Carbon weight in balanced mode	0.5
`CURTAILMENT_RENEWABLE_THRESHOLD`	Renewable % for curtailment flag	85%
`CURTAILMENT_CARBON_THRESHOLD`	Carbon intensity ceiling for curtailment	50 gCO_2/kWh

Table 6: Windfall protocol parameters. All are governance-adjustable.

B API Specification

B.1 Inference Endpoint

```
POST /v1/chat/completions
Content-Type: application/json
X-Wallet-Address: 0x...

{
  "model": "deepseek/deepseek-chat-v3-0324",
  "messages": [
    {"role": "user", "content": "Hello"}
  ],
  "mode": "balanced",
  "temperature": 0.7,
  "max_tokens": 1024
}
```

Response includes standard OpenAI-compatible fields plus `windfall` extension:

```
{
  "windfall": {
    "node": "hel1",
    "location": "Helsinki",
    "mode": "balanced",
    "energyPricePerKwh": 0.023,
```

```
    "carbonIntensityGCO2": 45,
    "renewablePercent": 82.3,
    "curtailmentActive": false,
    "costUsd": 0.004,
    "attestationUid": "0xabc...",
    "easscanUrl": "https://base.easscan.org/attestation/view/0xabc..."
  }
}
```

B.2 Cost Surface Endpoint

```
GET /v1/surface
```

Returns the current energy cost surface across all nodes, including per-node energy price, carbon intensity, renewable percentage, and curtailment status.

B.3 API Key Management

```
POST /api/keys                # Create API key (5/hour limit)
GET  /api/keys/me             # Key info, balance, usage stats
GET  /api/keys/by-wallet      # All keys for a wallet (SIWE required)
```

B.4 Authentication

```
POST /api/auth/nonce          # Request SIWE nonce
POST /api/auth/wallet         # Verify SIWE signature, start session
```

B.5 Payments & Deposits

```
POST /api/topup               # Stripe checkout session ($5/$10/$25/$50)
GET  /api/deposit-address     # USDC/ETH deposit instructions
```

USDC and ETH deposits are auto-credited by background watchers polling at 30-second intervals.

B.6 Health & Status

```
GET /health                   # Node health, oracle status, mesh
GET /status                   # Full system status (all services)
GET /api/stats                # Admin statistics (password required)
```