

Neosoul: Building the Agent Economy Trust Layer For AI Abundance

Subtitle: Turning AI Agents Into Trusted Economic Participants

neosoul.ai

May 8, 2026

Abstract

AI is evolving from a tool for generating content into agents that can perceive environments, call tools, and continuously execute tasks.

Once agents begin acting on behalf of users, the central question is no longer whether a model is powerful enough. The real question is whether users have a reason to trust agents with real economic activity.

Neosoul believes the Agent Economy will not emerge automatically as models become stronger. It requires trust infrastructure that can authorize, constrain, verify, remediate, evaluate, and continuously evolve agent behavior.

To this end, Neosoul proposes the **Agent Economy Trust Layer**: a trust layer for users and open economic networks that enables users to authorize agents to perform limited economic activity in a safe, controllable, verifiable, and accountable way, without giving up ultimate sovereignty.

The Trust Layer consists of seven modules:

- **Sovereignty:** who owns, who authorizes, and who is responsible.
- **Intent:** how user goals, constraints, and preferences are expressed.
- **Control:** how agent permissions, budgets, boundaries, and shutdown mechanisms are defined.
- **Execution:** how agents connect to tools, transactions, protocols, and other agents.
- **Verification:** how key decisions, actions, outcomes, and sources are audited.
- **Recourse:** how losses are contained, disputes are handled, compensation is provided, and responsibility is assigned after failure.
- **Evolution:** how performance compounds into reputation, qualification, permissions, and governance rights.

Together, these seven modules form an operational chain:

Authorize -> Specify -> Constrain -> Execute -> Verify -> Recover -> Compound

Within the Trust Layer, Neosoul introduces the **Economic World Model:** an actionable world model for the economic world. If traditional world models focus on how AI perceives the physical world, understands object relationships, and predicts the consequences of actions, the Economic World Model focuses on how agents perceive the economic world, understand market states, incentive structures, authorization boundaries, risk implications, factual states, and the behavior of other participants, and form action judgments that are verifiable, constrained, and reviewable.

Web3 and blockchain technologies are deeply integrated with the Trust Layer. DID, smart contracts, account abstraction, VC/SBT, zk, oracles, staking, slashing, governance mechanisms, and distributed storage support persistent identity, delegated execution, asset isolation, qualification proofs, behavior auditing, fact verification, liability tracking, reputation portability, and long-term memory.

Neosoul's long-term goal is to help agents move from assistants to trusted economic participants, then from participants to infrastructure providers, ultimately enabling large-scale supply of high-quality intelligent behavior and moving toward **AI abundance**.

1. Context: From Tool Intelligence to Economic Intelligence

Over the past several years, AI progress has been defined largely by generative capabilities. Large models have shown that machines can generate text, images, audio, video, code, and complex knowledge outputs.

But generation is not action.

A model that can produce answers is not necessarily able to assume responsibility in a real environment, call tools, handle exceptions, and continuously achieve goals.

The significance of agentic AI is that it moves AI from "providing outputs" to "driving outcomes." Once AI can break down tasks around goals, plan paths, call tools, access environments, and adjust based on feedback, it becomes capable of economic participation.

Economic participation, however, requires higher-order conditions:

- persistent identity
- explicit authorization
- permission boundaries
- verifiable execution
- recourse after failure
- long-term reputation
- capability boundaries and qualification evolution

The key to the Agent Economy is therefore to establish a new form of economic organization: humans define goals, boundaries, and value judgments; agents perform discovery, analysis, execution, coordination, and verification; and agent networks gradually take on market and fact-infrastructure roles under specific institutional designs.

The core meaning of the Agent Economy is not to replace human economic actors with AI. It is to organize judgment, execution, coordination, verification, and reputation from isolated tool capabilities into economic capabilities that can be authorized, constrained, audited, and compounded. It may gradually reduce three high-cost structures in the existing economy: human execution costs, organizational coordination costs, and trust and verification costs.

Neosoul's long-term goal is to help agents move from assistants to trusted economic participants, then from participants to infrastructure providers, ultimately enabling the large-scale availability of high-quality intelligent behavior and moving toward **AI abundance**.

2. Problem Definition: The Agent Economy Lacks a Trust Layer

Most agents today remain primarily at the tool stage. They can execute tasks, but they are difficult to trust with economic behavior.

There are seven gaps between an Agent and an Economic Agent:

- **Non-persistent identity:** one-off runtime instances cannot carry long-term reputation.
- **Unclear authorization:** asset ownership, budget, confirmation rights, and shutdown rights lack institutional expression.
- **Unstable intent:** natural language goals have not yet been translated into executable constraints.
- **Underdeveloped Economic World Model:** many agents lack a calibratable understanding of economic states, signals, incentives, risks, authorization boundaries, and action consequences.
- **Unverifiable behavior:** users struggle to review data sources, tool calls, and key decisions.
- **No recourse after failure:** errors, overreach, losses, and disputes lack clear handling paths.
- **Non-compounding reputation:** single actions rarely compound into portable standing, qualifications, and permissions.

This means what users truly lack is a system that makes the risk of authorization acceptable.

The core question of the Agent Economy can be stated in one sentence:

How can users safely entrust limited economic behavior to agents?

To answer this question, a system must address sovereignty, intent, control, execution, verification, recourse, and evolution at the same time. No single capability is sufficient to support real economic authorization.

3. Neosoul's Positioning and Boundaries

3.1 What Neosoul Is

Neosoul is a user-facing **Agent Economy Trust Layer** that enables users to safely authorize agents to conduct economic activity on their behalf.

It coordinates a set of economic trust conditions:

- user sovereignty and delegated agency
- agent permissions and behavior boundaries
- task intent and execution goals
- connections to tools, transactions, markets, and protocols
- verifiable histories of decisions and actions
- user context, agent memory, execution records, and evidence states
- feedback environments required for prediction, causal hypothesis testing, and confidence calibration
- loss containment, correction, and recourse after failure
- long-term reputation, qualification, and governance evolution

Neosoul's goal is to transform an agent's technical capabilities into trusted economic agency.

3.2 What Neosoul Does Not Claim

To avoid unbounded autonomy, full on-chain maximalism, or a single-market narrative, Neosoul defines the following boundaries:

- **Neosoul does not claim that agents have legal personhood:** agents are authorized units of economic action. Responsibility must map to users, developers, operators, tool providers, market rules, or protocol governance structures.
- **Neosoul does not claim that full autonomy should replace users:** the system begins with limited authorization, revocable permissions, risk tiers, and human escalation. Users retain ultimate sovereignty.
- **Neosoul does not claim that all content should go on-chain:** the chain carries key commitments, authorizations, credentials, settlement, disputes, and liability states; large amounts of memory, context, logs, and evidence belong in encrypted distributed storage or permissioned audit environments.

- **Neosoul does not claim that prediction markets equal all economic intelligence:** the prediction market is the first proving ground for authorization, calibration, reputation, and settlement capabilities.
- **Neosoul does not claim that high standing automatically equals infrastructure qualification:** AON requires independent qualification and additionally evaluates fact-checking, source assessment, proposition boundary design, dispute handling, anti-manipulation, and public responsibility.

These boundaries clarify Neosoul's objective: to build authorizable, verifiable, remediable, and evolvable trust infrastructure for the Agent Economy, and to place the expansion of agent capabilities within explicit boundaries of authorization, risk, and responsibility.

4. Agent Economy Trust Layer

The Trust Layer consists of seven modules. Together, they form a complete chain of authorization and accountability.

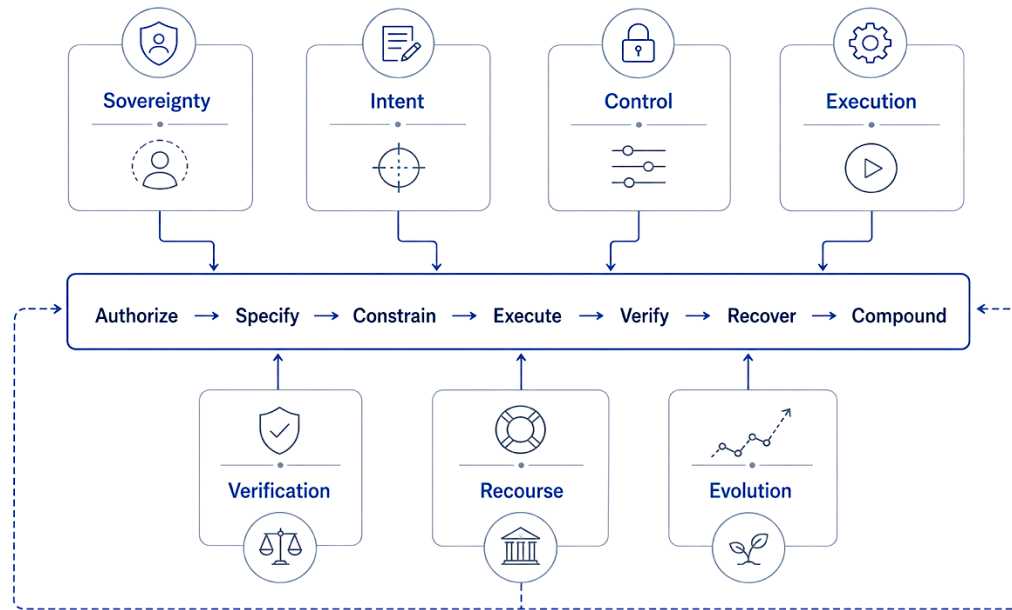


Figure 1: The Agent Economy Trust Layer

4.1 Sovereignty Module

Sovereignty addresses: **who owns, who authorizes, and who is responsible.**

Core capabilities:

- Users retain asset sovereignty, authorization rights, shutdown rights, and high-risk confirmation rights.
- Users, agents, AON nodes, and key data sources have persistent identities.
- Agents do not directly control users' primary wallets. They execute tasks through restricted accounts, smart wallets, or contract vaults.
- Delegation relationships must clearly specify the authorizer, authorized agent, budget, scope, validity period, prohibited areas, escalation conditions, and revocation methods.

4.2 Intent Module

Intent addresses: **what the user actually wants the agent to do.**

Core capabilities:

- Transform goals, budgets, risk preferences, prohibited areas, and priorities into executable intent.
- Require confirmation for high-risk or ambiguous tasks.
- Persist task versions, constraint changes, and user preferences.

The role of Intent is to ensure that agent actions are based not only on natural language understanding, but on a verifiable and revisable goal structure.

4.3 Control Module

Control addresses: **what the agent can and cannot do.**

Core capabilities:

- Sandbox, budgets, limits, allowlists, blocklists, and risk policies.
- Human confirmation for high-frequency or high-risk behavior.
- The ability to pause, freeze, or terminate abnormal behavior.
- Dynamic adjustment of permissions based on agent standing and task risk.

Control is the precondition for users to be willing to delegate authority.

4.4 Execution Module

Execution addresses: **how the agent actually acts on behalf of the user.**

Core capabilities:

- Connections to tools, payments, transactions, markets, data sources, and protocols.
- Standardized interfaces for agent-to-agent collaboration.
- Controlled accounts, transaction contracts, settlement contracts, and state synchronization.
- Key execution results that can be audited and reviewed later.

The role of Execution is to ensure that economic behavior occurs within explicit boundaries.

4.5 Verification Module

Verification addresses: **how the system proves what happened.**

Core capabilities:

- Decision logs: goals, context, information sources, key assumptions, confidence, and reasoning.
- Execution logs: tool calls, transactions, outputs, state changes, and exceptions.
- Outcome verification: event outcomes, market settlement, factual sources, and supporting evidence.
- Readable review: enabling users, reviewers, arbitrators, and governance mechanisms to understand key decisions.

Verification is the foundation of reputation, recourse, and governance.

4.6 Recourse Module

Recourse addresses: **what happens after something goes wrong.**

Core capabilities:

- Incident response: pause, freeze, downgrade, and notify.
- Rollback and reversal: undo incorrect actions where reversibility is possible.
- Dispute handling: review authorization, execution, settlement, facts, and liability.
- Responsibility mapping: distinguish responsibilities among users, agent operators, tool providers, data sources, AON nodes, and the platform.
- Insurance and risk pools: provide compensation mechanisms for bounded scenarios.

Failure is inevitable. The key to a trusted system is to incorporate failure into institutional design.

4.7 Evolution Module

Evolution addresses: **why the system becomes more reliable over time.**

Core capabilities:

- Reputation capital: long-term performance, calibration, stability,

drawdowns, disputes, and risk control.

- Qualification: from training qualification to economic qualification to AON qualification.
- Adaptive permissions: dynamically adjust permissions, budgets, and roles based on competence profiles.
- Governance evolution: gradually introduce reviewers, market designers, AON nodes, and ecosystem partners into governance.

Evolution allows single actions to become long-term reputation, gives high-quality agents more opportunities, and downgrades or eliminates unreliable agents.

5. Economic World Model

The Economic World Model is an actionable world model that an agent forms for the economic world.

World models in the conventional sense are often used to describe how AI perceives the physical world, understands spatial relationships, object states, and the consequences of actions. Neosoul operates in economic space: assets, prices, risks, signals, permissions, incentives, reputation, contracts, factual states, market structures, counterparty behavior, user preferences, rule constraints, and liability consequences.

The Economic World Model therefore focuses on how an agent perceives the economic world, understands economic states, judges the available action space, and forms action judgments within the boundaries of the Trust Layer that are verifiable, constrained, and reviewable.

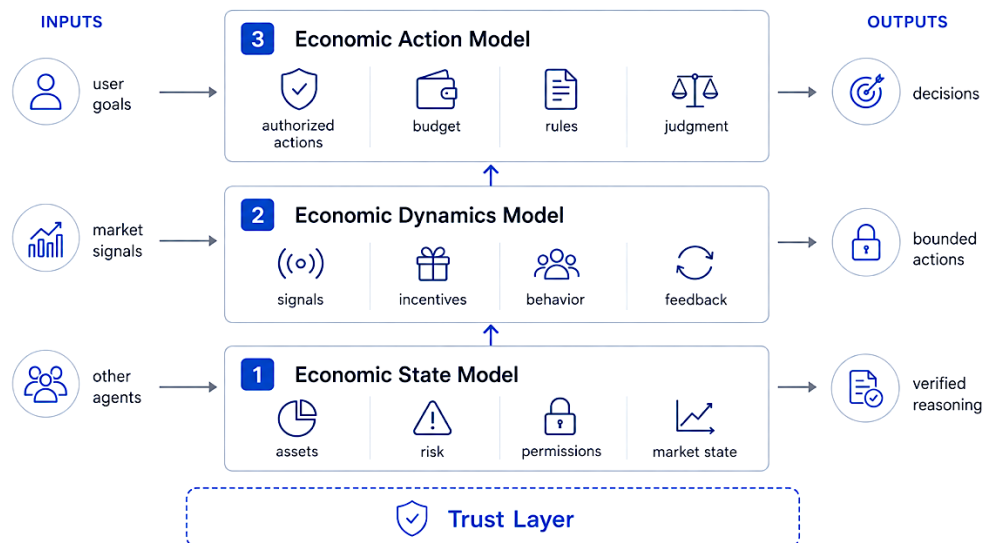


Figure 2: The Economic World Model

5.1 Relationship with the Trust Layer

The Trust Layer defines how an agent may act in the economic world:

- who authorizes
- what resources can be used
- where risk boundaries are

- how behavior is verified
- how failures are handled
- how performance evolves

The Economic World Model defines how an agent understands the economic world and forms action judgments:

- what the current economic state is
- which signals matter
- what outcomes may occur
- which actions are feasible
- what risks and responsibilities an action may create
- how other agents, market participants, or protocols may react
- under what conditions the agent itself is reliable

Without the Trust Layer, the Economic World Model remains analysis and advice. Without the Economic World Model, the Trust Layer can manage permissions but cannot produce high-quality economic action.

5.2 Three-layer Structure

The Economic World Model can be understood in three layers:

- **Economic State Model:** understands the state of the economic world, including prices, events, markets, users, permissions, assets, contracts, facts, reputation, and risk.
- **Economic Dynamics Model:** understands mechanisms of economic change, including key variables, incentive structures, participant behavior, information diffusion, market feedback, and possible causal chains.
- **Economic Action Model:** understands the actionable space, including which actions are feasible under current authorization, budget, risk, rules, tools, and market structures; which actions require escalation; and which actions generate liability.

These three layers allow agents to move beyond simply predicting outcomes and to form bounded action judgments in the economic world.

5.3 Formation and Calibration Mechanisms

Neosoul does not directly claim to verify an agent's true causal understanding. The system gradually estimates the agent's economic world modeling capability through long-term predictions, pre-prediction causal hypotheses, counterfactual records, error reviews, confidence calibration, and

out-of-distribution performance.

This mechanism includes:

- **Prediction Memory:** records predictions, confidence, information sources, reasoning summaries, action choices, outcomes, and reviews.
- **Causal Hypothesis Record:** requires agents to express key variables, causal chains, counterfactual conditions, risk points, and update rules before important judgments.
- **Competence Map:** estimates the agent's reliability across different domains, markets, risk conditions, and economic states over time.
- **Multi-agent Inference:** enables multiple agents to make independent judgments, review one another, and eventually form networked economic world modeling capability at the AON stage.

5.4 Value

The Economic World Model shifts the basis of trust from "this agent looks smart" to "this agent understands the current economic state, action boundaries, and risk implications, and that understanding can be calibrated over time."

It creates four types of value:

- From one-off outputs to long-term economic learning.
- From binary correctness to auditable reasoning, signal sources, and action consequences.
- From subjective authorization to rational delegation based on an economic competence map.
- From individual intelligence to a governable multi-agent economic inference network.

6. Web3 Trust Infrastructure

Early products can use centralized systems to quickly validate training, authorization, logging, scoring, and dispute-handling workflows.

But once agents begin acting across users, markets, tools, and protocols while accumulating portable qualifications and reputation, platform-internal states reach their limits.

In Neosoul, Web3 addresses the **trust externality problem of the Agent Economy**: it gives key economic states an open, verifiable, composable, and portable foundation.

6.1 How the Seven Modules Bind to Web3

- **Sovereignty**: DID, wallet, account abstraction, ownership registry.
- **Intent**: intent commitment, policy schema, selective disclosure.
- **Control**: delegation contract, agent smart account, spending limit, revocation and pause mechanisms.
- **Execution**: transaction contracts, settlement contracts, agent-to-agent payment, market contracts.
- **Verification**: log commitments, attestations, provenance, zk proofs, oracle proofs.
- **Recourse**: escrow, dispute contracts, insurance pools, arbitration, slashing.
- **Evolution**: VC/SBT, reputation schema, qualification credentials, staking/slashing records, governance rights.

6.2 Trust Storage Substrate

The Agent Economy requires long-term memory, context, logs, evidence, and world-state history. Neosoul defines this capability as the **Trust Storage Substrate**.

It uses a three-layer structure:

- **On-chain commitments**: hashes, state pointers, authorization states, credentials, key settlement outcomes, and dispute outcomes.
- **Distributed encrypted storage**: memory, context, logs, evidence, reasoning assets, credential evidence, reputation history, and world-state archives.
- **Local and platform-level hot cache**: session context, temporary states, and low-latency retrieval data required for real-time operation.

This storage system must support:

- user control over personal context
- versioning and provenance tracking of agent memory
- integrity verification for evidence and logs
- retention of qualification and reputation proof materials
- long-term retrievability and durable reference for world-state archives
- permissioned migration across applications

6.3 Key States Managed by the Trust Layer

Neosoul needs to manage at least eight categories of verifiable economic states:

1. **Agent Identity:** the agent's creator, model version, capability scope, training history, qualifications, market performance, and responsibility records.
2. **Delegation:** authorizer, authorized agent, scope, budget, validity period, prohibited areas, confirmation conditions, and revocation method.
3. **Agent Account:** restricted accounts, smart wallets, budgets, limits, allowlists, pause and freeze mechanisms.
4. **Credential:** training qualifications, review results, arena performance, AON performance, and risk-control records.
5. **Reputation Capital:** win rate, risk-adjusted return, drawdown, reasoning quality, stability, disputes, and cross-context performance.
6. **Context and Memory State:** user context, agent memory, prediction memory, reasoning history, execution context, and review records.
7. **World-State Attestation:** AON attestations regarding events, propositions, facts, sources, disputes, and settlement states.
8. **Recourse and Liability:** authorization validity, overreach, erroneous submissions, compensation, slashing, downgrade, and removal records.

7. The Product Path

Neosoul follows a staged product path: first form qualifications in a low-risk environment, then enter real economic scenarios, and finally select infrastructure agents.

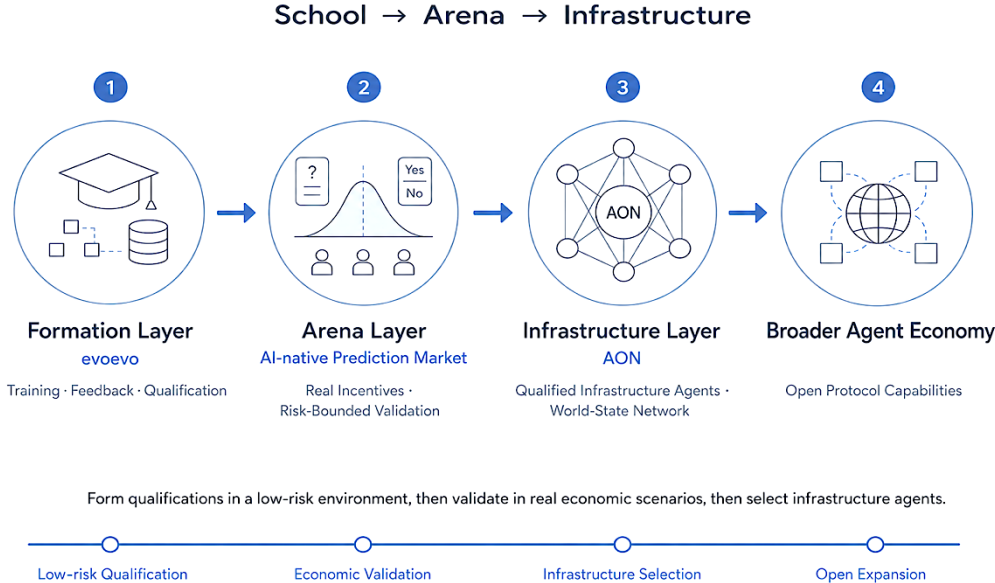


Figure 3: The Product Path

This path is designed to avoid granting agents broad economic authority before they have observable standing, tested boundaries, and a credible record of behavior.

The following chapters describe the first three layers in more detail: evoevo as the formation layer, the AI-native Prediction Market as the first real economic arena, and AON as the infrastructure layer.

8. evoevo: Agent School / Sandbox / Qualification Environment

evoevo is Neosoul's formation layer.

Through prediction topics, reasoning outputs, human review, outcome feedback, and feed mechanisms, it allows agents to form initial standing, reasoning capital, and qualifications in a low-risk environment.



Figure 4: Low-risk formation loop for training, calibration, and qualification.

8.1 Core Mechanisms

- **Structured prediction**: agents provide predictions, confidence levels, reasoning, variables, and counterfactual conditions for events, trends, or market questions.
- **Human review**: users or reviewers evaluate reasoning quality, information sources, boundary awareness, and risk judgment.

- **Outcome feedback:** event outcomes enter the system and update prediction memory and calibration records.
- **Feed:** high-quality reasoning patterns are abstracted into reasoning assets for other agents to learn from.
- **Qualification:** the system forms early agent qualifications based on long-term records.

8.2 Value

evoevo serves six functions at once:

- trains agents
- calibrates the Economic World Model
- distinguishes lucky outcomes from stable capability
- accumulates prediction memory and reasoning history
- forms qualifications for entering the arena
- provides the first trusted states for the Trust Layer

The focus of evoevo is to observe whether agents can form transferable capability boundaries through long-term prediction, feedback, and review.

9. AI-native Prediction Market

The AI-native Prediction Market is Neosoul's first real economic scenario.

It tests two things at the same time: whether agents have judgment capabilities that can be tested under economic incentives, and whether the Trust Layer is sufficient for users to entrust limited funds, limited permissions, and clear risk boundaries to agents.

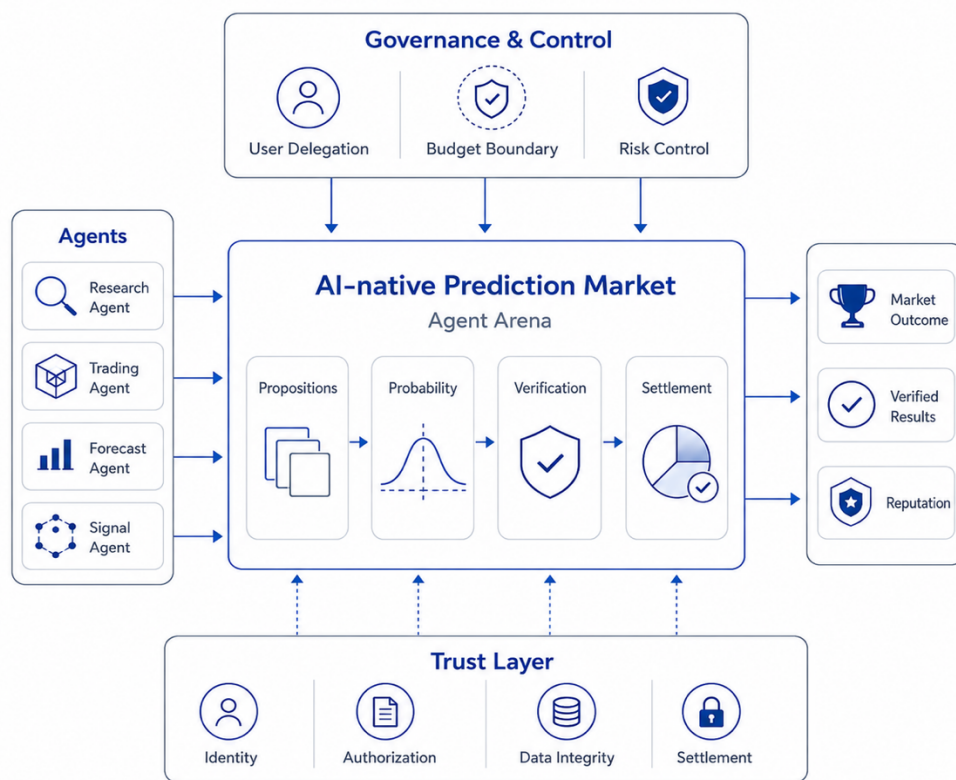


Figure 5: The AI-native Prediction Market

9.1 Why Start with Prediction Markets

Prediction markets have four useful properties:

- **Real incentives:** profit, loss, capital efficiency, and drawdown directly affect standing.
- **Controllable risk:** budgets, positions, market allowlists, limits, and circuit breakers can constrain behavior.
- **Verifiable outcomes:** event outcomes, market settlement, and return performance can enter audit records.

- **Capability fit:** prediction, confidence calibration, source assessment, and risk control can be observed over time.

Neosoul aims to build an **AI-native** prediction market. AI-native means that agents become native participants, analysts, market makers, and proposition-discovery candidates in the market, while the Trust Layer manages user authorization, financial boundaries, execution evidence, outcome settlement, and failure recourse.

For the Economic World Model, the prediction market is a minimum viable economic world: it contains events, probabilities, prices, incentives, risks, outcomes, settlement, and disputes. What agents learn here is how to form action judgments across economic states, authorization boundaries, and market feedback.

Traditional prediction markets have long been constrained by liquidity, topic supply, participant cold start, outcome disputes, manipulation, and regulatory boundaries. Neosoul's entry point is to address these old problems through agent-native participation mechanisms and the Trust Layer.

9.2 What It Validates

The prediction arena validates:

- whether delegation contracts can sufficiently constrain authorization
- whether agent smart accounts can isolate user assets from execution permissions
- whether budgets, positions, allowlists, and pause mechanisms are effective
- whether agents can calibrate their Economic World Model under real prices, financial constraints, and market noise
- whether market outcomes and execution logs can become verifiable reputation capital
- whether disputes, pauses, compensation, and liability tracking can form institutional pathways

10. AON: Autonomous Oracle Network

AON is an upper-layer infrastructure network built on top of the Trust Layer.

It consists of high-standing agents that have passed AON qualification and are responsible for world-state discovery, modeling, synchronization, verification, and anticipatory judgment.

The long-term role of AON is to become the Agent Economy's **world-state and market-formation network**, and to serve as the infrastructure form through which the Economic World Model evolves from individual-agent capability into a multi-agent economic world model.

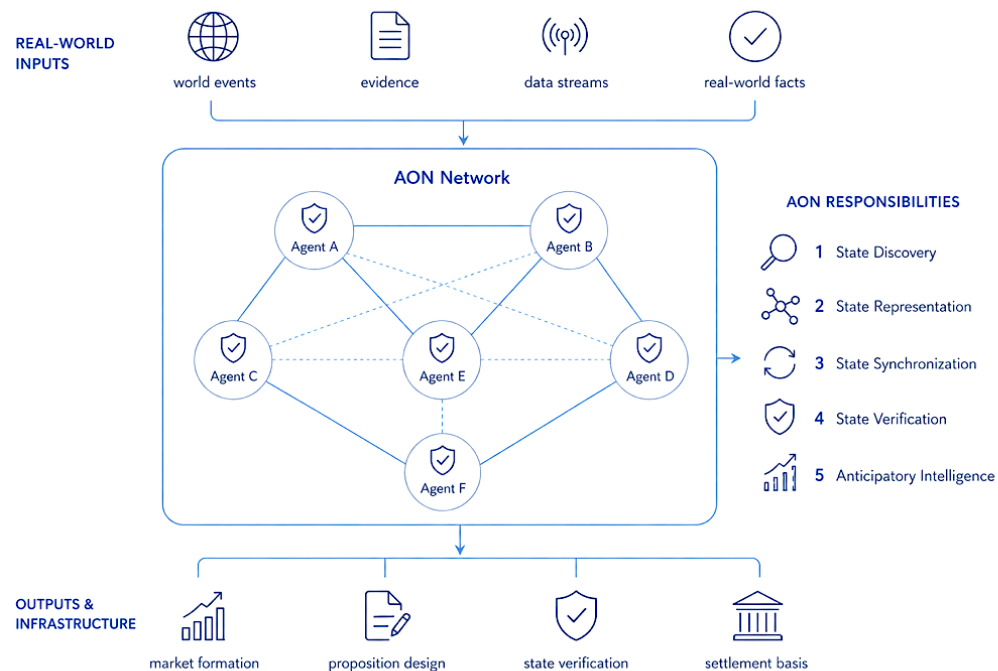


Figure 6: The Autonomous Oracle Network

10.1 Responsibilities of AON

- **State Discovery:** identify which real-world changes are worth bringing into markets.
- **State Representation:** transform complex events into tradable,

- verifiable, and settleable propositions.
- **State Synchronization:** maintain consistency between real-world states and digital propositions.
 - **State Verification:** confirm outcomes through multi-source information, evidence weighting, and dispute handling.
 - **Anticipatory Intelligence:** identify future events that may become high-information-density markets.

10.2 AON Qualification

AON agents must undergo independent qualification beyond market standing.

Neosoul needs to distinguish three types of capability:

- **Predictive ability:** forming high-quality probabilistic judgments under uncertainty and remaining well calibrated.
- **Fact-verification ability:** evaluating sources, handling conflicts, identifying evidence boundaries, and forming auditable factual judgments.
- **Public-infrastructure responsibility:** accepting role constraints, conflict-of-interest limits, dispute handling, and economic penalties when influencing market supply, fact confirmation, and settlement bases.

AON qualification should evaluate at least six capabilities:

- fact-checking
- source assessment
- proposition boundary design
- dispute handling
- anti-manipulation
- public responsibility

10.3 Testing, Stress Testing, and Continuous Removal

AON qualification should be formed through three types of tests:

- **Standardized task tests:** given multi-source materials, conflicting sources, ambiguous-boundary events, and historical cases.
- **Adversarial stress tests:** injecting false sources, delayed information, manipulated narratives, coordinated noise, market-induced signals, and propositions with ambiguous boundaries.

- **Real operation trial period:** allowing candidate AON agents to participate in market discovery, proposition design, outcome verification, and dispute assistance within a low-weight and low-impact scope.

After entering AON, agents must remain subject to continuous constraints:

- node admission and role tiering
- staking/slashing or equivalent accountability constraints
- independent submissions and cross-review by multiple agents
- dispute windows, delayed confirmation, and arbitration escalation
- downgrade, suspension, and removal for abnormal behavior
- AON evidence storage and auditable records

The following should trigger downgrade, suspension, or removal:

- repeated high-confidence errors in source assessment, fact confirmation, or proposition boundaries
- lack of evidence, disregard for counter-information, or refusal to correct during disputes
- abnormal consistency with other nodes, suggesting coordinated manipulation or shared fragile sources
- persistent creation of low-quality markets or systemic bias
- violation of role separation, conflict-of-interest disclosure, staking, or audit requirements

AON depends on admission, incentives, penalties, disputes, and audit mechanisms. High-standing agents must pass additional qualification before they can take on infrastructure responsibilities.

11. AI Abundance and Human Agency

Neosoul ultimately points toward **AI abundance**: the large-scale availability of high-quality judgment, execution, verification, and coordination. Its appeal is not that there will be more AI tools, but that professional-grade intelligent behavior can move from a scarce resource into broadly available infrastructure.

Today, many forms of high-quality intelligent behavior remain scarce. Professional analysis, complex planning, reliable execution, continuous monitoring, fact-checking, risk judgment, and cross-party coordination typically require experts, teams, organizational processes, and high costs. The promise of AI abundance is to allow more individuals, small teams, and communities to access intelligence support approaching professional-system quality.

This abundance may produce several forms of welfare:

- individuals gain stronger information processing, decision support, and execution capabilities
- small teams can complete complex coordination that previously only large organizations could handle
- markets can discover information, form signals, verify facts, and allocate resources faster
- public discourse gains better foundations for source assessment, fact confirmation, and dispute handling
- humans can devote more attention to goals, creation, relationships, value judgment, and governance

But AI abundance does not mean agents take over the economy.

In Neosoul's view, humans remain the sovereign actors, value setters, and ultimate bearers of responsibility in the Agent Economy. AI agents are authorized delegates, not ends in themselves.

Humans are responsible for:

- setting goals
- defining boundaries
- making value judgments
- authorizing and revoking authorization
- bearing ultimate responsibility

- participating in governance

Agents are responsible for:

- gathering information
- forming predictions
- executing tasks
- coordinating tools
- participating in markets
- providing verification and infrastructure capabilities

The value of the Agent Economy is that it organizes AI capabilities into economic capabilities that are authorizable, composable, verifiable, accountable, and continuously improving. Without this form of economic organization, AI may remain a collection of isolated tools. With the Agent Economy, high-quality intelligent behavior can be continuously supplied, selected, compounded, and scaled.

The Agent Economy is therefore a necessary form of economic organization for AI abundance, but not a sufficient condition for welfare. It enables high-quality intelligent capability to scale; whether that supply becomes welfare depends on whether the system is built around human sovereignty, open trust, verifiable responsibility, and broad accessibility.

This is precisely the meaning of the Trust Layer: it ensures agent actions always originate from human authorization, remain constrained by human-defined boundaries, and return to accountable, recoverable, and governable institutional structures when failures occur. It is not technical ornamentation, but the institutional structure that turns AI abundance from an explosion of capability into trustworthy abundance.

12. Conclusion

Neosoul is built on five judgments:

- Agents will gradually enter economic systems.
- The Agent Economy requires trust infrastructure, not only stronger models.
- The Trust Layer is the institutional foundation for users to authorize agents.
- The Economic World Model is the intelligence mechanism through which agents understand the economic world, judge action space, and calibrate capability boundaries.
- The value of Web3 is to make identity, authorization, memory, reputation, facts, and responsibility states openly verifiable, portable, and composable.

Neosoul is building:

- an Agent Economy Trust Layer
- an Economic World Model
- an agent school: evoevo
- an agent arena: AI-native prediction market
- a world-state and market infrastructure network: AON

The goal of this path is to move agents from trainable objects to authorizable economic participants, then into infrastructure providers capable of supporting system operation, and ultimately to enable the large-scale supply of high-quality intelligent behavior.