

# HeyElsa —The Crypto Agent Layer ( Litepaper )

---

## Elsa Overview

TL;DR

Elsa is the **crypto agent layer** that turns **intent → action**.

Users (and partners) issue goals; a **multi-agent system** plans, validates, and **executes across chains** safely, autonomously, and at scale.

---

## ▼ Overview (What is Elsa?)

Elsa makes DeFi feel like typing a message.

Three surfaces, one engine:

- **Copilot (B2C):** Chat to swap, bridge, stake, hedge, farm, set automations.
- **Widget & SDK (B2B):** Embed “trade with AI” inside any wallet, dApp, or content app.
- **AgentOS:** Build/host specialized agents on a shared **Agent-to-Agent (A2A)** coordination bus.

---

## ▼ What this litepaper covers (Preview)

- **Vision → Autonomous DeFi:** From on-demand execution to **self-driving** portfolio actions (take-profit, hedging, rebalancing, APY hopping).
- **Multi-LLM Orchestration:** Route tasks by complexity/latency to the right model; verified, grounded outputs.
- **A2A Communication Layer:** Specialized agents coordinate, bid for work, and share context.
- **Execution & Safety:** Simulation, guardrails, MEV-aware routing, receipts, and proofs.

- **Cognitive Cache:** Privacy-preserving memory for personalization without leakage.
- **Build on Elsa:** Agent OS, SDKs, embeddable widget; host your own agents and monetize.

---

## ▼ Technical Architecture (Layered)

**Bottom → Top**

### 1. Data & Grounding

On-chain state (nodes/indexers), oracles, curated KB; freshness checks & anomaly detection.

### 2. Execution Layer

Pre-audited scripts, simulation, route selection (DEX/bridge meta-routers), **MEV-safe submission**, idempotency, receipts, telemetry.

### 3. Agent Layer

Composable specialists (Swap, Bridge, Yield, Risk/Hedge, Perps, NFT, Sniper, Alerts). First-party + **hosted third-party** agents with scoped permissions & SLAs.

### 4. A2A Bus (Coordination)

Pub/sub messaging, contract-net bidding, plan aggregation, reputation, fairness scheduling.

### 5. Orchestration (Multi-LLM + Planner)

Task classification → model routing (small/mid/large). Strategy synthesis, constraint solving, fallback/hedged inference.

### 6. Intent Layer (NL Interface)

Parsing, slot-filling, constraint capture, policy hints, multi-turn context & disambiguation.

### 7. Safety & Policy

Simulation gates, allow/deny lists, max-notional limits, geo/KYC hooks, **verified inference (zkTLS/MPC-TLS)**.

### 8. Cognitive Cache

TEE-backed private memory of preferences, behaviors, outcomes; accelerates planning & personalization.

## 9. Observability

Traces, metrics, audit logs, anomaly detection, circuit breakers.

---

# How the three surfaces map onto the stack

- **Copilot:** Intent → Orchestration → **A2A** → Agents → Execution  
(Safety, Cache, Grounding applied throughout)
- **Widget & SDK:** Partners attach at Intent/Orchestration with their context; reuse the same **A2A**, Agents, Execution.
- **Agent OS:** Lives inside Agent/A2A layers; provides SDKs, sandboxing, manifests, reputation, and monetization for hosted agents.

---

# Layers, Policies & Safety

## Intent Layer

**Role:** Natural-language interface that converts user input into machine-readable intents with constraints.

### Responsibilities

- Parse NL (text/voice) into structured `Intent{verb, objects[], params{}, constraints{}, riskHint}.`
- Slot filling (amounts, assets, chains, time horizons, targets) with uncertainty scores.
- Constraint capture: risk limits (max slippage, LTV caps), compliance (geo/KYC), budgets, deadlines.
- Context injection from Cognitive Cache (prefs, past approvals, favorite chains/pools).

### Inputs → Outputs

- In: NL prompt, wallet context, session policy, cache profile.
- Out: Validated intent JSON + confidence; missing-slot prompts if needed.

### Key Components

- NER + typed entity registry (tokens, chains, protocols).

- Canonical schema & versioning (e.g., `v3.intent.swap`, `v2.intent.yield.optimize`).
- Guardrails (regex, allow-lists, refusal rules for unsafe verbs).

### Failure Modes / Handling

- Low confidence → ask clarifying question.
- Ambiguous asset/chain → propose top-3 disambiguation.
- Out-of-policy verb → safe decline with rationale.

---

## Orchestration Layer

**Role:** Decides *how* to satisfy an intent: routes to the right model(s), synthesizes a plan, assigns agents, and sets up A2A messaging.

### Responsibilities

- Task classification (complexity, domain, required latency, risk class).
- Multi-LLM routing (small/mid/large or toolformer) with cost/latency budget.
- Plan synthesis: decompose into steps with pre/post-conditions and success metrics.
- Agent assignment & capability matching; create A2A contracts for cooperation.

### Inputs → Outputs

- In: Intent JSON, cache profile, live data pointers (oracles/indexers).
- Out: `Plan{steps[], dependencies[], SLAs, fallbackPaths[], evalFns[]}` + agent roster.

### Key Components

- Classifier features: **task complexity, risk class R0–R3, latency budget, chain/protocol access, wallet approvals needed, MEV exposure, statefulness needs.**
- Routing policy (see section below).
- Planner (STRIPS-like) with constraint solver (fees, gas, liquidity, bridge risk).

### Failure Modes / Handling

- No feasible plan → return alternatives (cheaper chain, different pool).

- SLA breach predicted → downgrade model, simplify plan, or defer.

---

## A2A Bus (Agent-to-Agent Communication)

**Role:** Contract-Net style coordination for multi-agent work; pub/sub plus scheduling and reputation.

### Responsibilities

- Matchmaking: broadcast sub-tasks; agents bid/commit with price/SLA.
- Reputation & scheduling: track success, latency, accuracy; prioritize reliable agents.
- Message semantics: idempotent commands, retries, exactly-once task receipts.

### Key Components

- Topics: `swap.quote`, `bridge.route`, `stake.position`, `risk.hedge`, `perps.open`, `alerts.set`.
- Message types: `rfq`, `proposal`, `award`, `work.start`, `work.done`, `work.fail`.
- Safety: per-task scopes, least-privilege creds, time-boxed leases.

---

## Agent Layer

**Role:** Composable specialists that execute parts of a plan.

### Core Agents

- **Swap Agent:** DEX quoting, split routes, slippage control.
- **Bridge Agent:** Route selection (risk/fee/time), proof checks, stuck-tx recovery.
- **Yield Agent:** APY scan, IL/risk modeling, auto-compound, vault migrations.
- **Risk/Hedge Agent:** Perps/options, stop-loss/TP, delta hedging.
- **Perps/NFT/Sniper/Alerts:** Venue-specific execution, floor sweeps, limit/cancel/IOC, event watchers.

### Interfaces

- Capability manifests (supported chains, tokens, max notional, expected latency).
- Deterministic simulation hooks for pre-trade checks.

- Telemetry contract: step outcome, gas used, PnL, confidence, anomalies.

---

## Execution Layer

**Role:** Turns approved plans into on-chain transactions safely.

### Responsibilities

- Pre-audited scripts per protocol (param-bounded).
- Transaction simulation (fork or RPC sim) with revert reason capture.
- Meta-routing: 0x/CoW/UniswapX/1inch; bridge aggregators; batching.
- MEV safety: private relays, back-run protection, non-arb bundle rules.
- Settlement & receipt collation; retries with nonces/gas bumping.

### Outputs

- `ExecutionReport{txids[], status, receipts[], realizedSlippage, fees, proofs}` .

---

## Cognitive Cache

**Role:** Private memory for speed and personalization.

### Contents

- Preferences (risk, fee sensitivity, favored chains/pools).
- Behavioral features (time-of-day activity, asset affinities).
- Allowances & trusted protocols.
- Summarized wallet history (positions, cost basis, PnL bands).

### Controls

- TTLs & decay, user export/delete, per-scope consent, encryption at rest.

### Benefits

- Fewer clarifications, better defaults, faster routing.

---

## Safety Layer

**Role:** Proves what we know, enforces what we allow.

### Grounding

- Live on-chain state (indexers/full nodes), oracles (prices/liquidity), curated KB (audits, risk advisories).
- Every recommendation/check passes `GroundingEngine.validate()`.

## Policy Engine

- Geo/KYC, protocol allow-lists, notional/risk caps, leverage ceilings, counterparty filters.

## Verified Inference

- zkTLS / MPC-TLS proof adapters for “this insight came from these sources at time T” without exposing raw logs.

## Simulation & Dry-Run

- Pre-execution sims with guard conditions and stop-if rules.

---

# Observability

**Role:** Make intents→actions traceable and operable.

## What's tracked

- Traces with span IDs from intent parse → plan → A2A → execution → receipts.
- SLAs/SLOs per step (p50/p95 latency, success rate, cost).
- Audit logs for all policy decisions and escalations.
- Anomaly detection (stuck bridge, abnormal slippage, oracle drift).

## Surfacing

- Operator dashboards, partner webhooks, user-visible receipts.

---

# Classifier Features (expanded)

- **Task complexity:** single-step vs multi-leg, cross-chain, approvals required.
- **Risk class (R0–R3):** read-only → low-risk swap → leveraged/perps → protocol-novel/high-risk.
- **Latency budget:** sub-second (quotes) to minutes (bridges) vs deferred (rebalances).

- **Chain access:** wallet approvals, token allowances, supported venues/bridges.
- **Market context:** volatility, liquidity depth, MEV risk bands.
- **User profile:** fee sensitivity, risk tolerance, preferred venues.

---

## Routing Policy (multi-LLM)

- **Fast / transactional** (balances, prices, parameter extraction) → **small finetunes** (low latency, tool use).
- **Analytical** (yield comparisons, pathfinding, risk scoring) → **mid models** with tool-augmented RAG.
- **Strategic / autonomous** (multi-leg plans, portfolio mgmt, hedging) → **large models** with planning + constraint solving.
- **Cost/latency governors:** per-request budgets; degrade gracefully under load.

---

## Grounding Pipeline

- **Fetch:** on-chain (nodes/indexers), oracles (price/liquidity), curated KB (audits, risk lists).
- **Validate:** cross-source consistency checks, freshness thresholds, quorum rules.
- **Attach:** include `groundingRef` in every recommendation/action.
- **Prove:** optional zkTLS/MPC-TLS proof artifact for partner/user verification.

---

## Verified Inference

- **What:** Cryptographic proofs that a model's answer was derived from stated inputs and logs.
- **How:** Proxy TLS via zk/MPC adapter; store digest; emit verifiable receipt with the response.
- **Where used:** High-stakes advice, partner compliance, dispute resolution.

---

## Fallbacks & Hedging

- **Model timeouts:** reroute to smaller model; summarize; request user confirmation if confidence < threshold.
- **Execution hedging:** secondary routes/venues; split orders; time-slice in volatile markets.
- **Dual-run critical plans:** plan via large model, cross-check via mid model; reconcile conflicts.

---

## Example Intent → Action (concise)

1. **User:** "Bridge 2 ETH to Solana and stake for safest yield."
2. **Intent Layer:** `intent.yield.optimize` with `amount=2 ETH`, `dstChain=Solana`, `risk=low`.
3. **Orchestration:** classify `R2`, route mid/large model → plan: {swap? no, bridge via Wormhole, stake Solana pool A}, SLAs.
4. **A2A Bus:** RFQs to Bridge Agent + Yield Agent; award best proposals.
5. **Execution:** simulate bridge + stake; private submit; collect receipts.
6. **Safety:** ground rates/APY; enforce policy caps; attach proof.
7. **Observability:** log trace; show user receipts + realized APY baseline.
8. **Cache:** update prefs (Solana OK, low-risk yield).

---

## Schemas (abridged)

### Intent (v3)

```
{
  "verb": "yield.optimize",
  "objects": [{"asset": "ETH", "amount": "2"}],
  "params": {"dstChain": "solana", "risk": "low"},
  "constraints": {"maxSlippageBps": 50, "deadlineSec": 900},
  "profileRef": "cache://user/123",
  "policyScope": "retail.low",
  "version": "v3"
}
```

### Plan (v2)

```
{
  "steps": [
    {"id": "s1", "type": "bridge", "route": "wormhole", "from": "ethereum", "to": "solana"}, {
      "id": "s2", "type": "stake", "venue": "solend", "asset": "SOL", "policy": "lowRisk"
    }
  ],
  "dependencies": [[{"s1", "s2"}]],
  "sla": {"p95LatencyMs": 30000},
  "fallbacks": [{"on": "bridgeFail", "route": "portal"}]
}
```

## SLAs / SLOs (suggested)

- Intent parse p95 < 400ms; plan synth p95 < 1.5s (analytical), < 4s (strategic).
- Quote freshness < 3s; simulation success > 99%; execution success > 98% (with retries).
- Max realized slippage ≤ quoted + policy buffer (e.g., +15 bps).
- A2A task award latency p95 < 800ms; agent failure rate < 1% per 1k tasks.

## Policy Defaults (retail tier)

- Max leverage 3×; LTV ≤ 60%; single-venue exposure ≤ 40% per action.
- Bridge preference: audited routes; fallback only with user confirm (R3).
- New protocol cool-down: require audit flag or manual review.

If you want, I can convert this into a Notion page export (as Markdown with toggles) or wire it into your existing data room structure.

```
flowchart LR
  I[Intent + Context] --> C[Complexity Classifier]
  C -->|Low| M1[Small Finetune]
  C -->|Medium| M2[Mid Model]
  C -->|High| M3[Large Model]
```

M1 → G[Grounding Checks]

M2 → G

M3 → G

G → P[Plan + Constraints]

## Build once; get execution + safety + distribution.

- **SDKs:** TypeScript/Python — intent schemas, on-chain connectors, simulation, policy hooks.
- **Lifecycle:** register → advertise capabilities → receive tasks → emit plan fragments/quotes/tx payloads → execute → report outcomes.
- **Sandbox:** per-agent runtime, quotas, scoped permissions (least privilege), secrets isolation.
- **Scheduling:** capability matching + live load + reputation.
- **A2A messaging:** pub/sub topics ( `quote` , `plan` , `risk` , `exec` , `status` ) with signed envelopes.

## Agent Manifest (example)

```
name: yield-agent
version: 1.2.0
capabilities:
  - chain: [base, ethereum, solana]
  - actions: [scan_yield, allocate, compound, exit]
sla: { p95_latency_ms: 1000, success_rate: 0.995 }
cost_model: { pricing: per-plan + per-exec }
permissions: { scopes: [simulation.read, protocols.write:yield, quotes.read] }
```

**Why:** multi-specialist coordination for complex workflows.

- **Protocol**
  - **Envelope:** signed JSON ( `msg_id` , `parent_id` , `topic` , `payload` , `ttl` , `sig` ).
  - **Transport:** NATS/Kafka-class pub/sub (durable streams), secure WS for interactive co-planning.
  - **Contract Net:** broadcast CFP; agents **bid**; coordinator selects on price/latency/reputation.

- **Consensus:** weighted voting on conflicting plans (reputation  $\times$  recency  $\times$  scope-fit).
- **Idempotency:** `(intent_id, step_id)`.
- **Trust & Reputation**
  - Metrics: accuracy, timeliness, slippage vs quote, failures, user feedback.
  - Decay to favor recency; quarantine/slash malicious or failing agents.

```
sequenceDiagram
Coordinator->>Agents: * Call for Proposals (intent, constraints)
Agents->>Coordinator: Quotes (cost/latency/risk)
Coordinator->>Agents: Award + Plan fragments
Agents->>Execution: Submit executable steps
Execution->>Coordinator: Receipts + Telemetry
Coordinator->>Agents: Reputation update
```

- **Pathfinding:** 0x / CoWSwap / UniswapX / 1inch meta-routing; bridge selector; order-splitting.
- **Simulation:** dry-run with pool state; bounds checks (minOut, gas ceiling, deadlines).
- **MEV:** private relays/bundles; price impact caps; back-run detection.
- **Accounts:** AA/ERC-4337; session keys; MPC/HW wallet support.
- **Risk:** dynamic slippage; conditional stops; liquidation buffers.
- **Receipts:** tx hash, logs, realized slippage, fees, route, quotes vs fills.
- **Idempotency:** replay-safe by `(user_id, intent_id, step_id)`.

```
flowchart LR
Plan --> S1[Simulate On-chain State]
S1 --> S2[MEV Risk Check]
S2 --> S3[Limits: Slippage, Gas, Notional]
S3 -->|Pass| TX[Submit Tx via Private Route]
S3 -->|Fail| ERR[Abort + Explain]
```

- **Content:** prompt history, wallet behaviors (chains, protocols, risk), strategy outcomes.
- **Store:** vector DB (embeddings) + key/value features.
- **Use:** disambiguation, defaults, constraint inference, few-shot exemplars.
- **Privacy:** TEE-backed; per-user scope; opt-in research sharing; differential privacy on aggregates.

**Excalidraw brief:** circular cache feeding router with “known preferences”, “risk profile”, “recent strategies”.

- **Sources:** on-chain nodes/indexers; oracles (price/TVL); curated KB (audits/safety).
- **Validators:** cross-source reconciliation, freshness, anomaly detection.
- **Policy engine:** denylist, chain allow-lists, max notional, KYC/geo hooks.
- **Proofs:** zkTLS/MPC-TLS adapters to prove gateway authenticity.
- **Explainers:** every block/reject returns a human-readable reason.
- **Modes:** execute now; schedule; **autopilot** (bounded by policies).
- **Features:** swaps, bridging, staking, perps (roadmap), copy trading, prediction markets, yield optimizer, sniping/limits, NFT trade & limit.
- **Personalization:** risk bands, chain prefs, fee sensitivity, favorites.
- **Outcomes:** ~\$100M total volume; 30M+ prompts; higher retention via context.

## Widget (drop-in)

- Embed: `<iframe src="https://app.heylesa.ai/widget?partner=XYZ&theme=dark" />`
- Prefill context (asset/chain/flow)
- Callbacks: `onQuote` , `onTx` , `onError` , `onComplete`

## SDK (agents & flows)

- Intent schema, plan API, simulation, submitTx, receipts
- Policy templates, sandbox scopes, quota mgmt
- Theme/tone overrides

```

import { Elsa } from "@elsa/sdk";

const elsa = new Elsa({ apiKey: process.env.ELSA_KEY });

const intent = {
  type: "yield.optimize",
  amount: "500",
  asset: "USDC",
  constraints: { chain: "base", maxSlippageBps: 30 }
};

const plan = await elsa.plan(intent);
await elsa.simulate(plan.id);
const receipt = await elsa.execute(plan.id);

```

- **Verification:** conformance tests (correctness/latency), replay suite, adversarial sims.
- **Runtime:** CPU/mem quotas, topic rate limits, cost ceilings.
- **Monetization:** per-plan/per-exec pricing (USDC/\$ELSA), rev share on partner volume.
- **Metrics:** intent→plan→exec funnel, p50/p95, slippage distribution, failure taxonomy, per-agent SLOs.
- **Tracing:** distributed trace IDs across router/agents/execution.
- **Alerts:** anomaly detection (slippage spikes, oracle drift), circuit breakers.
- **Audit:** immutable logs for compliance & partners.
- **Wallets:** MPC + HW; ERC-4337 AA; session keys with ceilings.
- **Exec:** allow-list bridges/routers; private flow; expiring quotes; gas caps; reorg-aware confirmations.
- **Inference proofs:** zkTLS/MPC-TLS verifier endpoints.
- **Data privacy:** encryption at rest; TEE for cognitive cache; opt-in telemetry.
- **Router:** p95 < 300ms (low/med), < 900ms (high).
- **Plan→Simulate:** p95 < 1.2s (single-chain), < 2.5s (cross-chain).

- **Quote accuracy:**  $\geq 99.0\%$  within bounds.
- **Execution success:**  $\geq 99.2\%$  after retries.
- **Throughput:** 200–500 intents/sec (horizontally scalable).

More users → richer prompts & outcomes → better finetunes → higher plan/exec quality → more partner conversion → more users.

flowchart LR

```

U[Users & Partners] --> D[Prompts + Outcomes Data]
D --> T[Finetuning + Policy Updates]
T --> Q[Higher Plan/Exec Quality]
Q --> A[Adoption & Volume]
A --> U

```

### Retail: "Bridge 2 ETH to Solana and stake best yield"

sequenceDiagram

```

User->>Copilot: "Bridge 2 ETH to Solana and stake best yield"
Copilot->>Router: classify + route (med/high)
Router->>Planner: multi-step strategy
Planner->>A2A: call for proposals (bridge,yield)
BridgeAgent-->Planner: quotes + routes
YieldAgent-->Planner: APY + risk + lockup
Planner->>Safety: simulate + policy + oracle checks
Safety-->Planner: OK w/ constraints
Planner->>Execution: submit (private relay)
Execution-->Chains: txs
Chains-->Execution: receipts
Execution-->Copilot: confirmations + proofs + explainer

```

### iFrame Widget

1. Paste snippet, set partner key, theme/tokens/chains.
2. Subscribe to callbacks for analytics/incentives.

### SDK

1. Install, create intents, plan/simulate/execute, display receipts.
2. Apply policy templates (max notional, allow-listed routers, geo limits).

## Host Your Agent

1. Publish manifest, pass conformance tests, get sandbox scopes.
2. Receive A2A tasks; earn per plan/exec.

```
POST /v1/intents
{
  "type": "swap",
  "user": "0x...",
  "params": {"sell": "ETH", "buy": "USDC", "amount": "1.0", "chain": "base"},
  "constraints": {"maxSlippageBps": 30, "deadlineSec": 90}
}
```

```
GET /v1/plans/{intentId}
# returns steps, quotes, risk, agents, policy gates
```

```
POST /v1/execute/{planId}
# returns tx receipts, route, realized slippage, gas
```

```
mutation {
  createIntent(input: {
    type: YIELD_ALLOCATE,
    asset: "USDC",
    amount: "500",
    chain: BASE
  }) { id status }
}
```

- **System Layers:** rectangle bands for Intent, Orchestration, Agents, Execution; sidecar “Safety”; circle “Cognitive Cache”.
- **A2A Market:** central “Coordinator” node; multiple Agents bid; arrows for CFP/quotes/award.
- **Safety Gates:** funnel: Plan → Sim → Policy → MEV → Submit.
- **Flywheel:** circular loop of Users → Data → Finetune → Quality → Adoption.

---

## Problem

Crypto is powerful but unusable at scale.

- Too many chains, tools, and workflows
- High cognitive load for basic actions
- Manual execution leads to poor outcomes
- Automation exists, but isn't accessible

Most users don't need more protocols — they need **better execution**.

---

## Solution

Elsa introduces an **AI-first execution layer**.

Users express intent.

Elsa executes it.

- Natural language → onchain actions
- Optimized routing across chains and protocols
- AI agents manage trading, yield, and risk
- Automation without sacrificing control

Elsa abstracts complexity while preserving transparency.

---

## Product Stack

### B2C: Elsa Copilot

- Swap, bridge, stake, borrow
- Perpetual trading and copy trading
- Yield optimization and portfolio automation
- Smart alerts and agent-driven execution

## B2B: Elsa Infrastructure

- Widget and iframe integrations
- SDKs for apps and wallets
- Agent execution APIs
- Revenue-sharing distribution model

---

## Business Model

Elsa is **usage-driven**, not speculation-driven.

Revenue sources:

- Transaction fees on execution
- AI agent commissions
- B2B SDK and widget licensing
- Inference and execution fees

As execution volume grows, revenue scales naturally.

---

## Token Overview

Parameter	Value
Token	\$ELSA
Total Supply	1,000,000,000
Decimals	18
Network	Base

\$ELSA is a **utility token**, powering execution, inference, and access.

---

## Utility of \$ELSA

\$ELSA is the economic backbone of Elsa.

## Fee Discounts

- Reduced execution fees for holders and stakers

## Feature Access

- Unlock premium agents (copy trading, sniping, cognitive cache)

## Gas Abstraction

- Gas-free transactions within Elsa using \$ELSA

## Inference & Execution

- AI model inference
- Onchain agent execution
- B2B SDK and widget usage

Utility scales with real platform usage.

---

## Token Allocation & Vesting

Category	Allocation	Vesting
Team	7%	12m cliff, 24m linear
Foundation	34.490%	20% TGE, 10m cliff, 24m linear
Community	40%	20% TGE, 48m linear
Pre-Seed	1.4%	12m cliff, 24m linear
Seed	9.11%	12m cliff, 24m linear
Liquidity	8%	100% at TGE

Community-first allocation with long-term alignment.

---

## Roadmap Snapshot

### Up to Q2 2025 (Completed)

- Public beta live
- Core DeFi Copilot launched

- \$1M+ daily volume
- 300K+ MAUs, 700K+ signups
- Monetization enabled

---

## Q3–Q4 2025

- Perpetuals via Hyperliquid
- Copy trading and sniping agents
- Yield vaults and automation
- Mobile apps (iOS & Android)
- Multi-agent orchestration
- \$ELSA Token Generation Event
- Token utility activation

---

## 2026

- Fully autonomous portfolio management
- Tokenized crypto indices
- Agent Builder SDK
- Automated strategy marketplace
- Cross-chain execution kits
- B2B scale across wallets and apps

---

## Design Philosophy

- **Execution > speculation**
- **Utility before governance**
- **Demand-backed token economics**
- **AI that acts, not just advises**

Elsa is built for scale, automation, and real usage.

---