TruCore

ATF Primer

Infrastructure for Autonomous Finance

The Problem

Autonomous AI agents are executing financial transactions at machine speed with zero human oversight. Current infrastructure assumes a human approves every material decision, an assumption that collapses when agents operate independently.

Without an external enforcement boundary, agents can exceed spend limits, interact with unapproved protocols, and leave no auditable trail. The result is uncontrolled capital exposure and no accountability.

AI models hallucinate, drift, and behave unpredictably under novel conditions. On-chain environments compound the risk: MEV bots, sandwich attacks, and front-runners actively exploit unprotected transactions. Capital preservation requires hard constraints that cannot be softened at runtime.

The Model

ATF enforces a four-stage deterministic pipeline on every agent transaction:

- Policy: declarative rules evaluated against the agent's intent before any execution begins.
- Permit: a scoped, time-bound authorization token granting minimal execution rights. Permits expire automatically and cannot be escalated.
- Validate: pre-flight simulation and constraint verification ensure the transaction conforms to policy before touching the chain. Fail-closed by default.
- Receipt: a cryptographic receipt captures every evaluation, approval, rejection, and settlement event for tamper-evident post-trade audit.

V1 Scope

V1 targets Solana with nine venue integrations across three categories:

- DEX: Jupiter, Orca, Raydium. Swap enforcement with slippage bounds and minimum-out checks.
- Lending: Solend, Marginfi, and Kamino (feature-gated). Collateral ratio and liquidation safeguards.
- Perps: Drift v2, Mango v4, Hyperliquid (feature-gated). Leverage caps, market allowlists, and notional limits.

Hard Invariants

Non-negotiable constraints enforced on every transaction. These cannot be bypassed, overridden, or weakened at runtime.

- Spend cap: max value per transaction and per rolling time window.

- Protocol allowlist: only pre-approved programs may be invoked. Jupiter, Orca, Raydium (swaps), Solend, Marginfi (lending), and Kamino (feature-gated). Perps adapters feature-gated, off by default.
- Slippage max: price deviation hard-capped with enforced minimum output.
- Cooldown period: minimum interval between high-risk actions.
- Permit TTL + nonce: permits expire fast and carry single-use nonces.
- Domain separation: each permit is scoped to a specific environment. Cross-domain reuse is invalid.

Threat Model

ATF is designed to mitigate the following categories of risk:

- Unbounded execution: agent submits transactions outside approved parameters, draining capital.
- Protocol drift: agent interacts with unapproved or compromised contracts.
- Slippage exploitation: adverse fills and MEV extraction erode portfolio value.
- Authorization creep: over-permissioned agents accumulate access beyond original scope.
- Audit opacity: no verifiable trail of what was checked, approved, or rejected.
- Adversary and MEV exploitation: sandwich attacks and front-runners extract value from unprotected transactions.

What Design Partners Get

TruCore is onboarding a small cohort of design partners for V1.

- Early API access with direct engineering support for integration.
- Influence on policy primitives, SDK design, and protocol coverage.
- Priority onboarding when ATF moves to general availability.

How to Engage

If you are building autonomous agents that execute financial transactions on Solana, apply as a design partner. We are looking for teams with live agents, real transaction volume, and a concrete need for execution guardrails.

This primer is intentionally narrow. TruCore is an umbrella for more products.

Apply: trucore.xyz/#waitlist