

Typed-Effect Kernels for Reliable Agent Systems

K-12, G-9, and M-7 · working preprint

Ryan Hunter, with AI research collaborators

August 2025 · public source edition September 2026

Abstract

Language-model “agents” have outgrown ad-hoc prompt chains. Production deployments demand typed contracts, auditability, policy enforcement, and cross-organization interaction. We introduce K-12, a typed-effect episode kernel that captures the environment and laws for one bounded run; G-9, a graph kernel for typed coordination (routing, scheduling, durability); and M-7, a market/exchange kernel for cross-organization capability leases, metering, and settlement. The trio is realized in the Effect TS calculus (Context Tags + Layers), enabling formal reasoning and hot-swappable back-ends. We provide normative laws, operational semantics, evaluation rubric against LangGraph-style workflows, and a conformance harness. This paper consolidates the specification, design rationale, and mappings to implementation.

This is an archival preprint. The later specifications and formalization work refine parts of it; read the current research essay for the operating argument.

Version: 0.1.2 (Spec Freeze) • Date: 2025-08-08 Effect Generic Order: A-E-R (Effect<Output, Error, Requirements>)

New in v0.1.2. We add canonical JSON schemas, end-to-end idempotency/“exactly-once” semantics, explicit OTel mappings, a minimal monitor grammar, structured capability/ABAC, Memory/ Ledger envelopes with erasure/tamper-evidence, and transport bindings for λ streams.

1 Introduction

Agent frameworks have proliferated, yet most blur three distinct concerns:

1. Local semantics of one bounded run (what we call an *Episode*).
2. Coordination of many runs across components and teams.
3. Economics/governance across organizational boundaries.

We separate these into three kernels:

- K-12 (Episode Kernel): Purity of control (δ/λ), explicit effect slots (Tools, Memory, Quota, Policy, Explain, Locale), and streaming outputs.
- G-9 (Graph Kernel): Typed nodes/edges, Router/Scheduler/Ledger separation, guarded cy-

cles, and metrics/monitor hooks.

- M-7 (Market Kernel): Directory, attestation, contracting, metering, settlement, dispute, reputation.

The architecture mirrors proven systems patterns: micro-kernel per run, orchestration fabric above, and market protocols outside trust boundaries.

2 Background & Goals

- Typed Effects. Effect TS models services as Context Tags and compositions as Layers, cleanly separating *ports* (interfaces) from *adapters* (back-ends).
- LLM Agents. We treat LLMs as tools inside Episodes: Observe/Orient may use embeddings and search; Decide may invoke a chat model; Act executes tools or sub-workflows.
- Goals. (i) Soundness (typed IO, replayable traces), (ii) Governance (quota, policy, locale), (iii) Scalability (routing/scheduling), (iv) Inter-org operation (markets).

3 K-12: The Episode Kernel

(See spec for details.) We define the stream chunk envelope (text/json/tool-call/result/error/system), cancellation/back-pressure, transport bindings (SSE/gRPC), error taxonomy, and right-to-erasure.

4 Persona: Sequencing Episodes

Persona is a tiny controller that decides when/how to run another Episode or delegate to a sub-graph; it lives above K-12. We formalize Transcript/Budget shapes and decision idempotency.

5 G-9: The Graph Kernel

Axioms on typing, cycles (guards), and routing idempotence. Options cover conditional edges, QoS & partitioning, lineage, human-in-the-loop, and GraphCall (trusted, in-org). We specify scheduling, parallel/join semantics, adapter nodes, and idempotency interactions across R and L.

6 M-7: The Market/Exchange Kernel (Alpha)

Directory, attestation, offers/contracts, metering, settlement, dispute, reputation. Laws to follow with early adopters.

7 Operational Semantics

Episode and Graph steps; idempotency keys propagate into L and R enabling safe retries. Ledger is tamper-evident; Memory uses LWW-element-set with tombstones for redaction.

8 Observability & Telemetry

IDs on spans/ledger; low-card metrics labels with exemplars keyed by trace/span; OTel mapping table enumerated in the spec.

9 Evaluation Rubric

Parity with LangGraph-style workflows; reliability under faults; governance checks; erasure correctness; observability correlation.

10 Security, Privacy, Locale

Least privilege; ABAC capabilities; Λ subset discipline; sandbox & egress control; RAG poisoning mitigations; embedding privacy; rotation of V leases.

11 Related Work

LangGraph, Airflow/Temporal, ReAct/OODA, Effect/ZIO calculus.

12 Limitations & Future Work

M-7 laws; formal monitors; TLA+ models; reference adapters.

13 Conclusion

K-12, G-9, and M-7 provide a principled, composable foundation for reliable agent systems.

Appendices

See spec/ documents for norms, schemas, and conformance suite.

References and context

- Shunyu Yao et al. 2023. [ReAct: Synergizing Reasoning and Acting in Language Models](#). *ICLR*.
- [LangGraph reference](#). LangChain.
- [Temporal workflow execution documentation](#). Temporal Technologies.

- [Managing Layers](#). Effect documentation.

Source edition notes

This edition preserves the August 2025 working preprint, with the unfilled release-commit field removed. The companion specifications appear in *The Kernel Trilogy: Source Specifications*. It is a research artifact, not a certification of every mathematical claim in the draft.