


The verifiable decision layer for the era of LLMs and agents.

KRML is a structured, naturalistic English that compiles, losslessly, into Answer Set Programming. **Flint & Steel** is the typed answer-set platform in active development underneath it. Schema projection turns one model into Protobuf, JSON Schema, OpenAPI, and SQL — every form a downstream consumer already speaks. Together: a structured decision platform whose authoring is readable, whose decisions are verifiable, and whose API is auto-projected.

THE PIPELINE



WHY THIS WORKS

Determinism on the answer

Same model and query, same answer. LLM variability stays in elicitation and narration, not in what's true.

Audit by construction

A literal proof chain in English. RAG systems show retrieved documents; this shows the inference itself.

Foundation-model independent

Better LLMs make the product better, with no retraining. The lab race is a tailwind, not a threat.

MARKET OPPORTUNITY

\$10–15B in existing categories — new ones forming now.

\$4–6B

Business Rules Management

Drools · IBM ODM · FICO Blaze · Pega — no generational refresh in fifteen years

Direct displacement

\$2–3B

Configuration & CPQ

Tacton · Configit · Oracle CPQ — strong in industrial, weak in regulated

Expand to new domains

\$2–4B

Scheduling & Planning

Per-industry point solutions; no shared foundation

Scheduling-as-policy

\$0.5–1B

Policy-as-code & authorization

OPA · Cedar · Oso · Permit.io — readability is the documented complaint

Own the readable tier

EMERGING — NO INCUMBENT

Verifiable agent governance

Every agent platform needs this within 24 months. The first credible product defines the category.

Structured layer for vertical AI

Hundreds of “AI for X” startups roll their own. We could be the platform they build on.

Regulator-facing decisions

EU AI Act, SR 11-7, SEC, state AI laws — the artifact that satisfies the regulator.

WHY NOW

- LLMs surfaced the gap: fluent reasoning that does not reproduce, cannot be audited step by step, and bends under pressure.
- EU AI Act, SR 11-7, SEC, state laws — regulators want inspectable logic and real audit trails.
- Every serious agent platform is hitting the action-governance wall now.
- Category-naming window for AI infrastructure: open now, closes by 2027.

WHY US

40 yrs of published KR research, Gelfond & Lifschitz to today
~16 peer-reviewed papers + 1 patent (AAAI, AAMAS, TPLP, AIJ)
86,000 lines of Rust, 1,900 tests, a working compiler
191 pp language reference, complete with worked examples

WHAT THIS IS NOT — BECAUSE OVERCLAIMING BURNS CATEGORIES

Not a replacement for LLMs — the structured decision layer above them. **Not a runtime verifier of free-form LLM behavior** — KRML verifies structured output, not prose. **Not the right tool for continuous-domain ML** — pair with the right tool there. **Not greenfield research** — the methodology is published, mature, production-ready. **Not automatically right about the law** — KRML guarantees the inference is correct given a faithful model; the modeling is human work, audited once.

“The foundation has been ready longer than the moment was. The moment is here.”

The ask: capital to build out the platform · early customers in regulated decision systems · distribution partners in agent platforms and vertical AI.