

ZeroTrain.Ai

Knowledge-Based Inference for Real-Time Decisioning

Forward Deployed Engineer Guide

**Building Deterministic Decision Systems for Agents, Applications, and Machine
Learning**

Prepared For:

Enterprise Clients, Architects, Regulators, Engineering Teams

By Leonard Gambrell

Executive Summary

Modern AI systems excel at generating language, identifying patterns, and making statistical predictions. However, many enterprise deployments eventually encounter the same challenge: the final business decision must be deterministic, explainable, auditable, and compliant.

Forward Deployed Engineers (FDEs) bridge the gap between cutting-edge AI capabilities and enterprise requirements. This playbook demonstrates how ZeroTrain.ai serves as a deterministic decision layer that integrates with existing AI architectures without replacing them.


1. The Enterprise AI Gap

LLMs generate text, ML predicts outcomes, applications automate workflows, and agents orchestrate tasks. Yet enterprises still require a deterministic layer that governs actions according to business policy, compliance, and governance.

2. Understanding the Decision Layer

ZeroTrain is not another LLM or ML model. It is a deterministic reasoning engine that converts explicit knowledge into repeatable decisions. Identical inputs and knowledge always produce identical outcomes.

3. Where ZeroTrain Fits



4. When to Use ZeroTrain

Use ZeroTrain where decisions matter more than predictions: finance, insurance, healthcare, manufacturing, cybersecurity, fraud detection, compliance, risk management, and enterprise workflows.

5. When NOT to Use ZeroTrain

Do not use ZeroTrain for creative generation such as chat, image generation, music, translation, or video generation. Allow generative AI to create information while ZeroTrain governs actions.

6. Discovery Questions

Ask: What decisions require manual approval? Where are business rules stored? Can decisions be audited? Can regulators reproduce yesterday's decision?

7. Common Integration Patterns (Examples)



8. Knowledge Engineering

Knowledge is authored directly instead of trained. Rules remain visible, reviewable, auditable, and independently versioned.

9. Validation

Validate knowledge correctness, expected outputs, deterministic repeatability, ONNX consistency, API consistency, and regression behavior.

10. Deployment



11. Operational Governance

Every deployment should answer: Why was this action taken?
Which rules fired? Who changed the knowledge? Can the decision be reproduced?

12. Best Practices

Keep knowledge modular. Version every artifact. Separate business policy from application code. Treat knowledge as software.

13. Frequently Asked Questions

- Does ZeroTrain replace LLMs? No. It governs deterministic decisions.
- Does it replace ML? No. ML predictions become inputs.
- Does it require GPUs? No.
- Does it require training? No.


14. Identifying ZeroTrain Opportunities

Look for long if/else chains, manual approvals, duplicated business rules, compliance gates, policy engines, spreadsheet-driven decisions, and inconsistent human decision making. These are strong indicators that a deterministic decision layer will add value.


Final Thoughts

Models generate possibilities. ZeroTrain governs decisions. Forward Deployed Engineers integrate deterministic reasoning into modern AI systems without replacing existing investments.

Agentic AI



Enterprise Applications



ML / Onnx Model Conversion

