

Beyond Compliance

The four technical debts actually blocking enterprise AI at scale — and why regulation isn't the one stalling your programme.

01 Security

Prompt injection, tool hijacking, model poisoning

02 Unit economics

Why pilot cost never predicts production cost

03 Data foundation

Fragmentation, retrieval decay, sync latency

04 Reliability

Silent failures in a probabilistic system

40%+

**of agentic AI projects will be cancelled
by 2027, Gartner forecasts.**

Cost, unclear business value, and inadequate risk controls are the leading reasons cited — not the model's capability.

SOURCE: GARTNER

The attack surface your security review can't see

- **Indirect prompt injection.** An agent reading an untrusted email or PDF can't reliably tell an attacker's hidden instruction from the user's real one.
- **Data exfiltration via inference.** Crafted inputs can pull proprietary or private data back out through the model's own output.
- **Supply chain poisoning.** Foundation weights, fine-tuning data, and open-source orchestrators are dependencies most security teams have never audited.

The pilot's cost model doesn't survive scale

- **Token spend compounds.** Multi-agent reasoning and recursive evaluation loops each multiply cost on their own.
- **The margin squeeze.** A workflow that clears ROI at 200 queries a day can lose money at 200,000 without caching and routing built in first.
- **Capacity is still scarce.** Dedicated accelerators mean capital committed months ahead, with fast depreciation and low off-peak utilisation.

60%

of AI projects will be abandoned through 2026 because the data was never made AI-ready.

Every layer above the data estate inherits its problems — a better model just produces confident wrong answers faster.

Everything above this inherits its problems

- **Fragmentation.** Knowledge trapped across shared drives and ticketing tools, with no standard metadata or access control.
- **Retrieval decay.** A vector store that answers well at 10,000 chunks degrades badly at 10 million without hybrid search and re-ranking.
- **Sync latency.** When an ERP record updates, "eventually" isn't good enough for the semantic index sitting downstream of it.

Failures that don't announce themselves

- **Silent hallucination.** A fabricated line item looks exactly like real output — nothing flags it automatically.
- **No continuous evaluation.** Providers update weights on their own schedule; most enterprises have no harness watching for regression.
- **Context decay.** Reasoning fidelity drops non-linearly over long sequences, well before the context window is technically full.

DOMAIN	FAILURE MODE	MITIGATION PATTERN
Security	Unsanitised tool execution	Gateway guardrails, least-privilege API scopes
FinOps	Unpredictable token spend	Semantic caching, tiered SLM/LLM routing
Data	Vector store pollution	Hybrid search, automated lineage sync
Reliability	Output drift, hallucination	Golden eval sets, LLM-as-judge, JSON schemas

Fix the architecture. The compliance conversation gets easier on its own.

Security boundary, then cost model, then data foundation, then evaluation harness — in that order, before production traffic, not after the incident report.



Terence Kok

AI & Innovation Strategist

[SAVE](#) · [COMMENT](#) · [FOLLOW](#)