Most RAG Pilots Are Framed as an LLM Problem. The Real Problem Is Governance.

A distinct discipline: secure, permission-aware, auditable knowledge systems — not just better prompts and vector stores.

SWIPE →

The Governed RAG 2.0 Stack

01 Design backwards from risk

02 Architect a policy-centric pipeline

03 Enforce permission-awareness at every layer

04 Build for auditability and security

4

Questions That Define the Governance Backbone

Scope. Risk classification. Obligations. Accountability model. Only once this is clear does it make sense to specify models or vector stores.

01

Every Retrieval Is a Fresh Authorisation Decision

Don't rely solely on front-end authentication. The policy engine filters retrieval results before the LLM ever sees them.

3

Typical failure modes: indexes built without row/column security, embeddings reused across sensitivity levels, and prompts that combine snippets from different tenants.

> **In high-stakes domains, the knowledge system must be an overlay on existing data controls, not a side channel.**

Terence Kok
Enterprise AI Strategist

RAG 1.0 vs. RAG 2.0

RAG 1.0	RAG 2.0
Choose a model, add a vector store	Policy engine filters before the LLM sees content
Index documents, ship a chatbot	Every layer carries identity + classification
Front-end authentication only	Immutable audit trail end-to-end

Treat the Model as an Untrusted Component.

Security has to extend beyond application security to prompt injection, retrieval poisoning, and indirect prompt attacks via retrieved content. Run high-sensitivity workloads in segregated environments; pass only the minimum necessary content to the model.

The Policy-Centric Pipeline

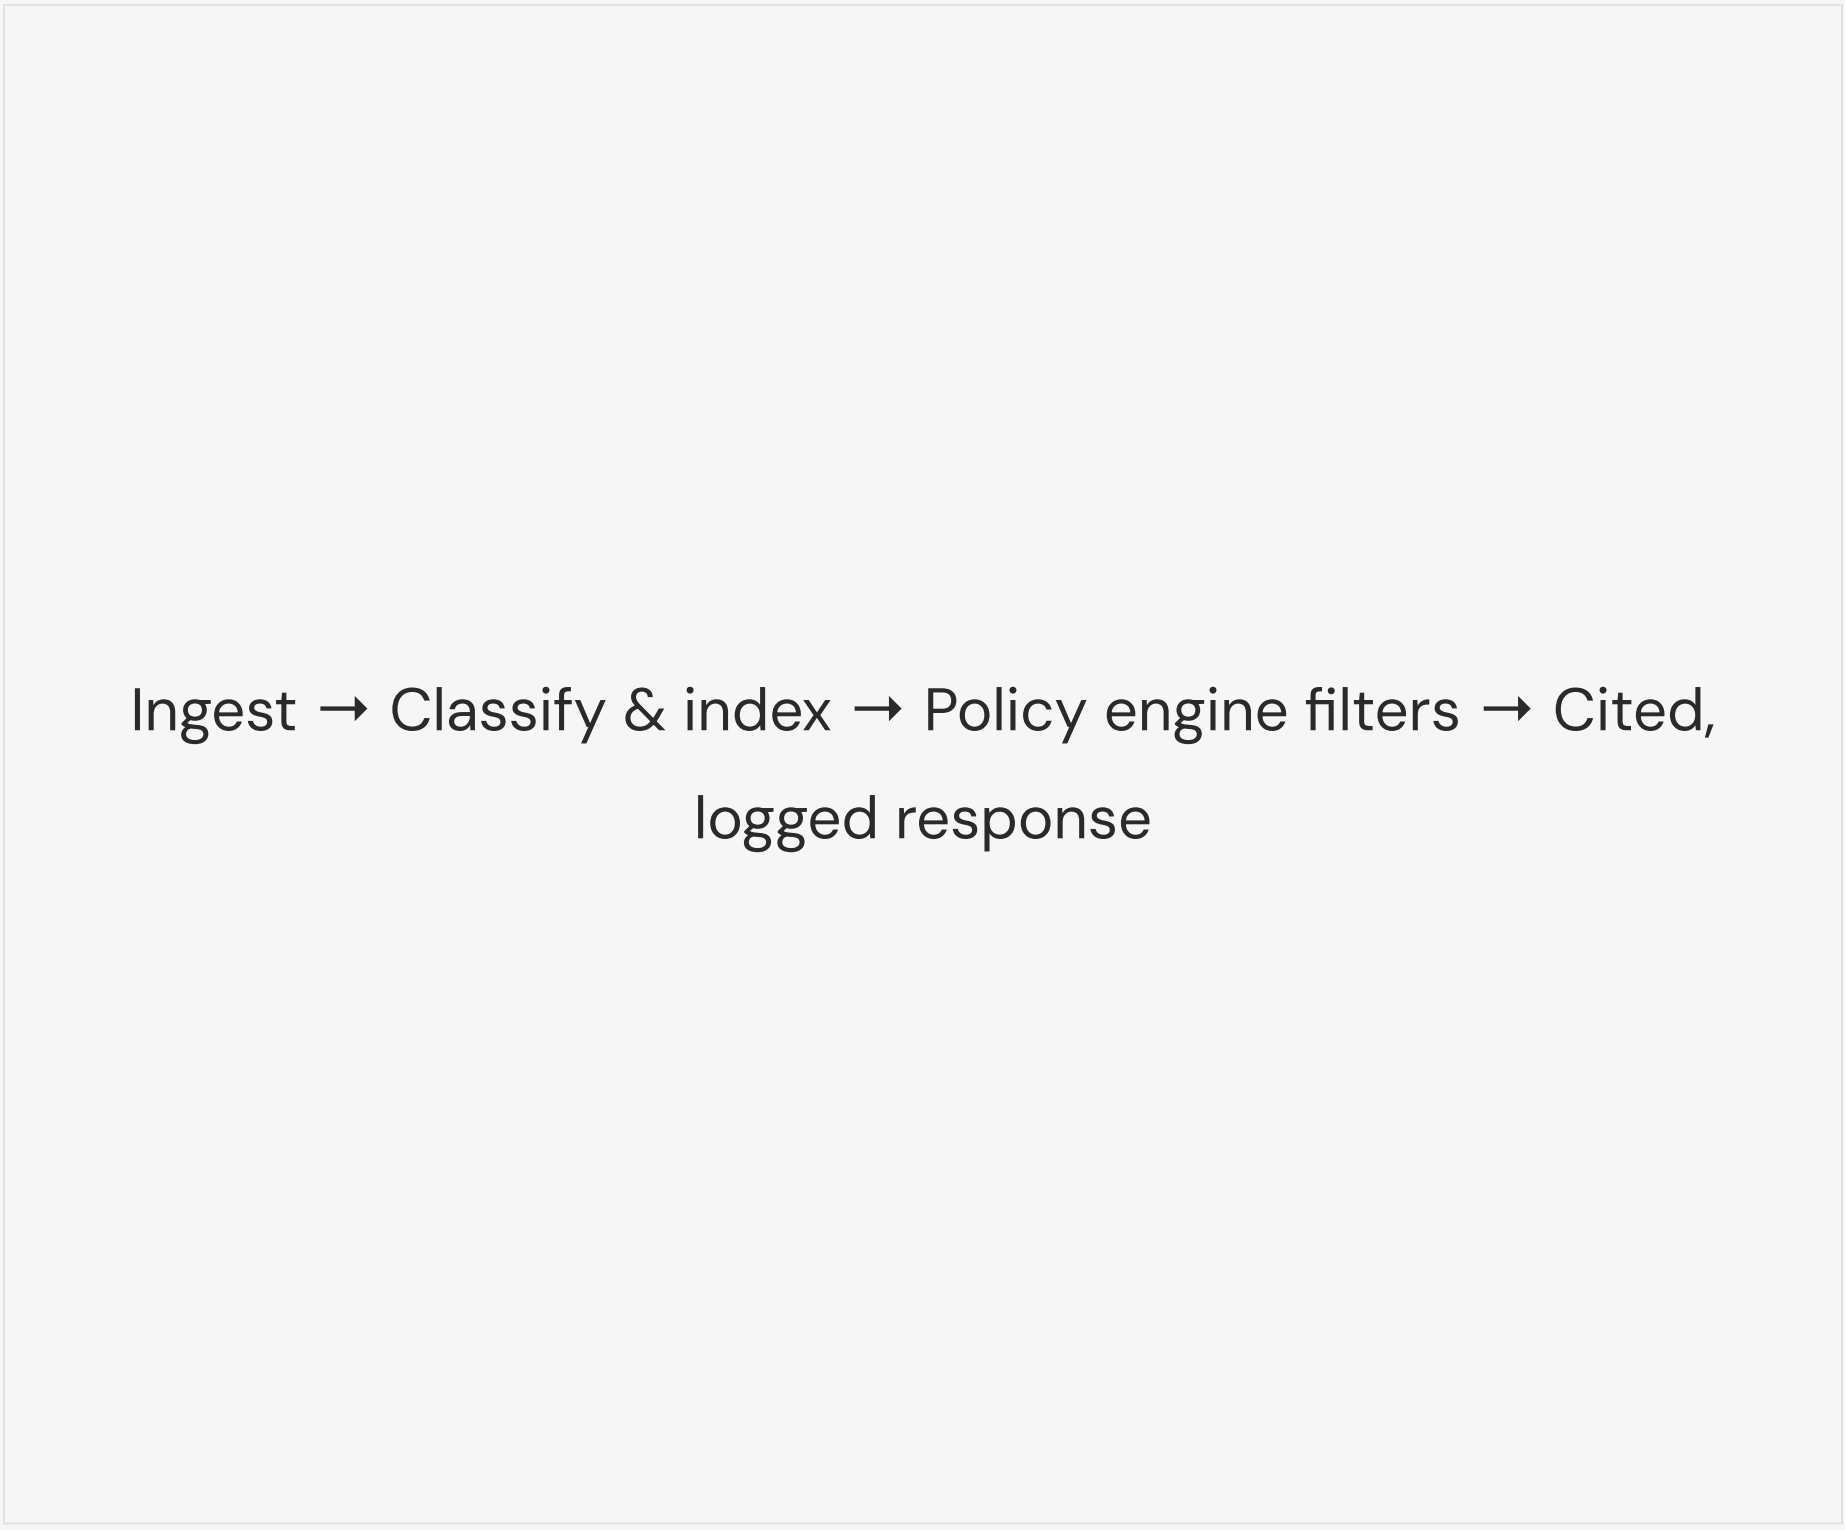

The full architecture, failure modes, and security checklist: terencekok.com/blog/how-build-governed-rag-20-systems-highstakes-use

Governed RAG Isn't Only About Avoiding Harm. It's a Strategic Capability.

How the organisation reasons over its own knowledge, end-to-end and under audit.

Terence Kok · Enterprise AI Strategist

SAVE · COMMENT · FOLLOW