

The Access Controls That **Disappear** When You Unify Data for AI.

Fixing fragmented data is step one. Rebuilding permissions at the AI layer is the step everyone skips.

Two Leaks, Two Fixes

01 Retrieval leakage vs. fine-tuning leakage

One serves the wrong document. The other memorises the fact into the model's weights.

02 The fine-tuning rule of thumb

Never train on data you wouldn't hand to every employee who can use the model.

03 Moving control from document to query

Every retrieval call needs to carry the requester's clearance, checked before the model sees anything.

Compare the Two Leaks

	RETRIEVAL LEAKAGE	FINE-TUNING LEAKAGE
Lives in	The vector database / document store	The model's own weights
Trigger	A query returns a document above clearance	The right prompt surfaces a memorised fact
Reversible?	Yes — fix the filter, leak stops	No — only retraining removes it



A model doesn't check anyone's job title before it answers.

If nobody rebuilt the wall, the model will happily explain it to whoever asked the right way.

BEFORE YOU FINE-TUNE, ASK

Would I hand this document to every employee with access to the model?

If the answer is no, keep it in a permission-aware retrieval layer instead — where access can still be revoked.

From Document to Query

OLD Access control lives on the file

A folder permission, set once, checked by a human opening it.

NEW Access control travels with every query

Row-level security and namespace filters check the requester before the model ever sees the data.

Classify. Gate. Redact. Log.

01 Classify before ingestion

Tier every document: safe to fine-tune, or retrieval-only and permission-gated.

02 Gate at query time

Filter results against the requester's clearance before the model reads anything.

03 Redact at ingestion

Mask sensitive fields when data is indexed, not after the model has already read it.

04 Log every retrieval

Record who asked, what was retrieved, and what the model returned.

Test It Like an Attacker Would.

Coax out restricted information with indirect prompts. Try it from a low-privilege account. Do this before launch, and on a recurring schedule after — configurations and datasets both drift over time.

The EU AI Act Expects a Paper Trail.

Clear data lineage and training-set documentation are required for higher-risk systems. An identity-tagged retrieval log is what makes that traceable — not just a policy on paper.

Fix the Data. Then Gate the Model.

Unifying data was never the finish line. Deciding what belongs in the model, and what stays behind a gate that can still be closed, is.



Terence Kok · AI & Innovation Strategist

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