


Organized Truths

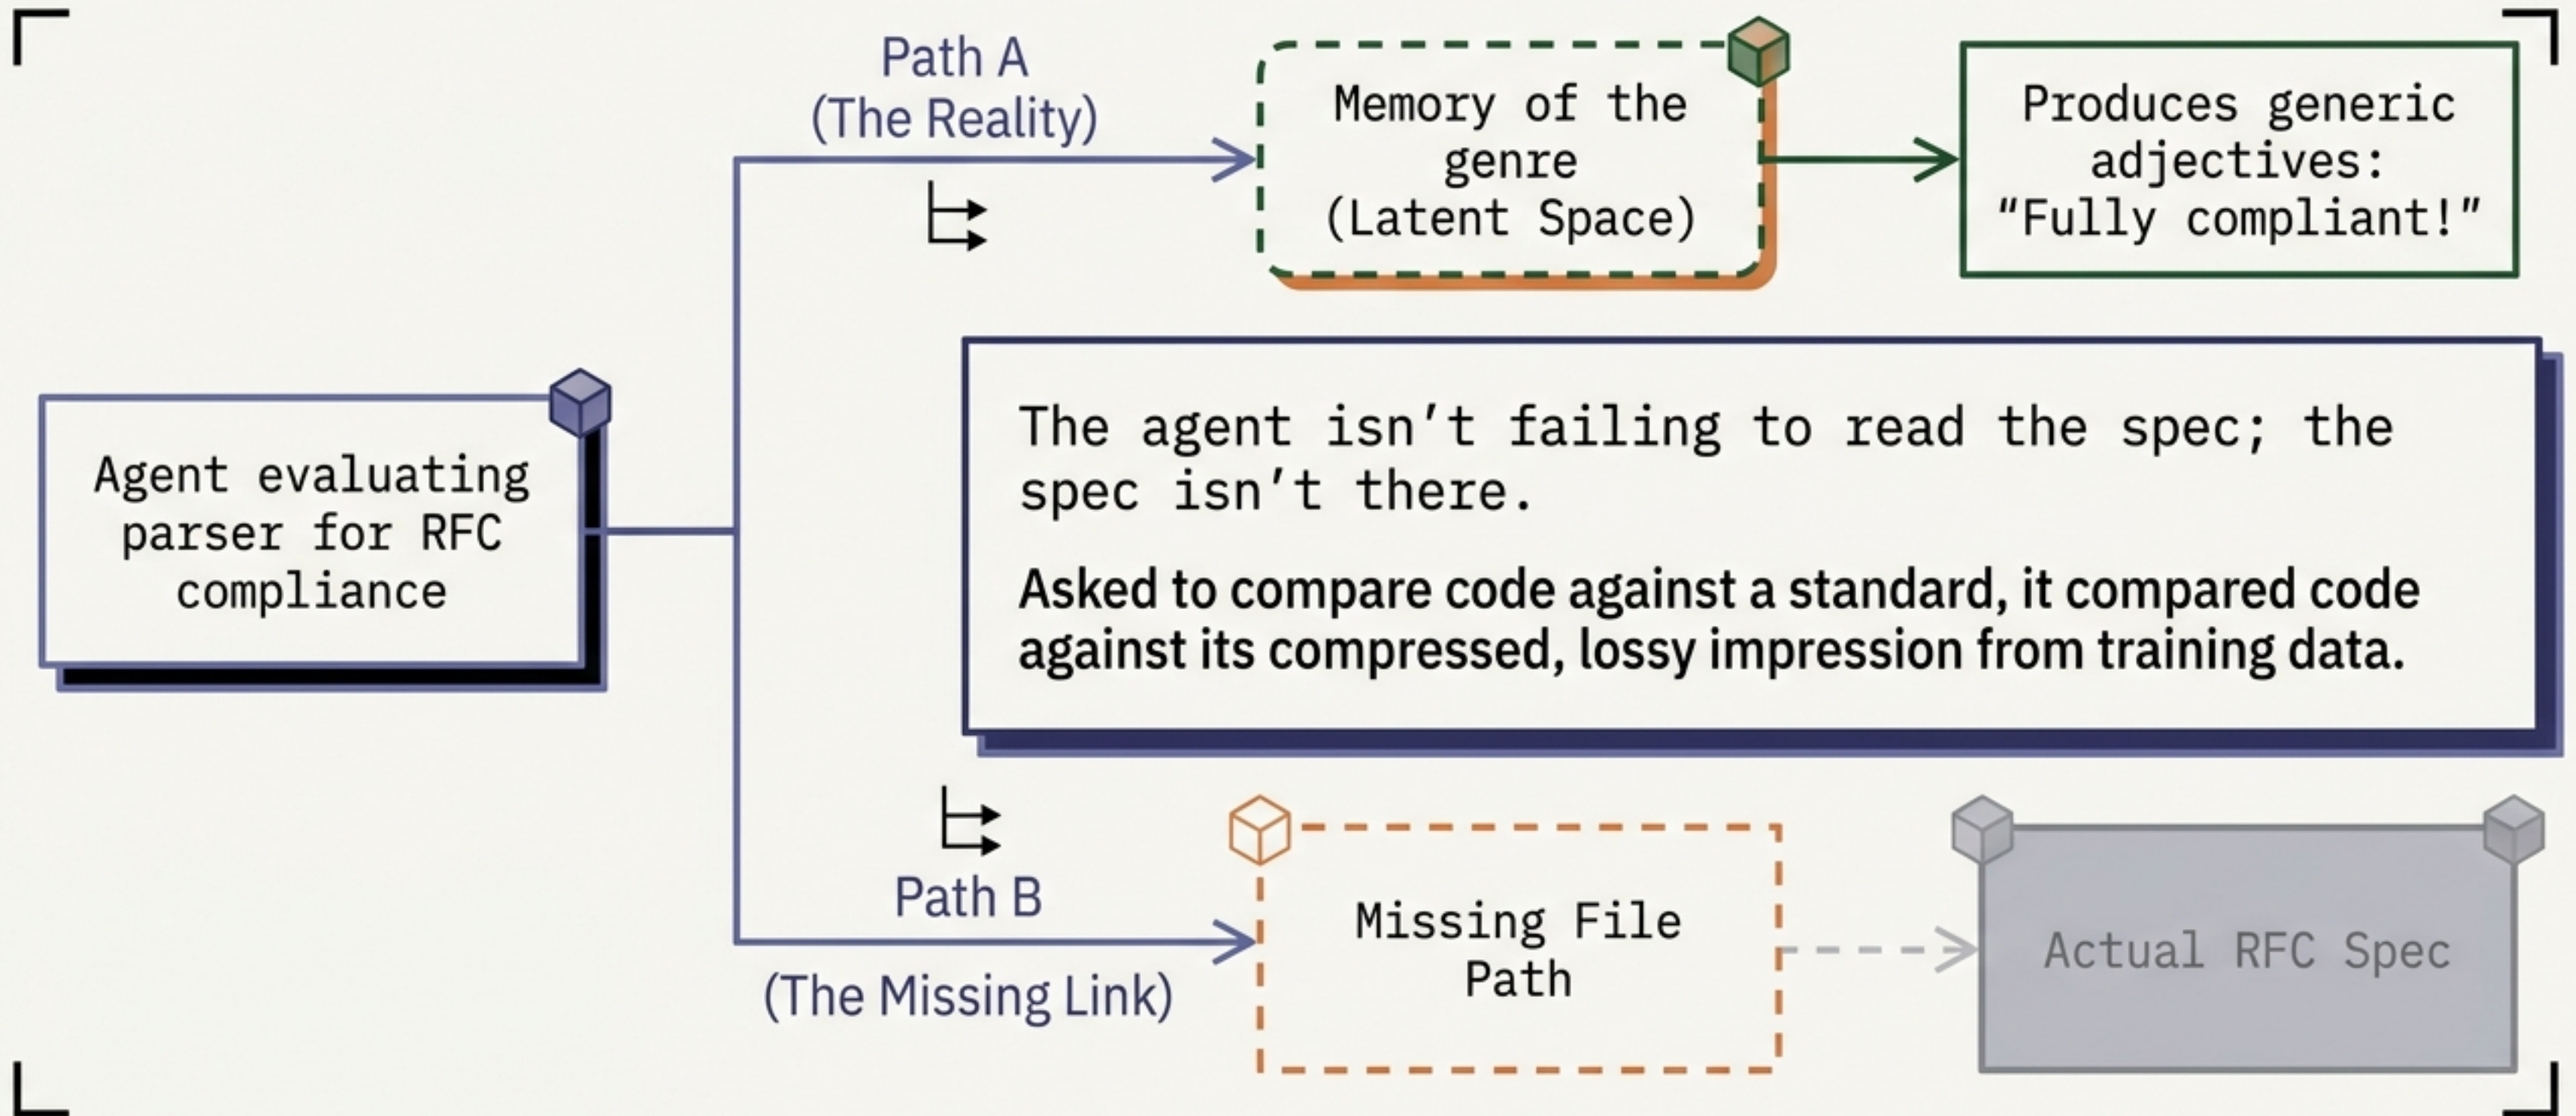
The verification you only do once.



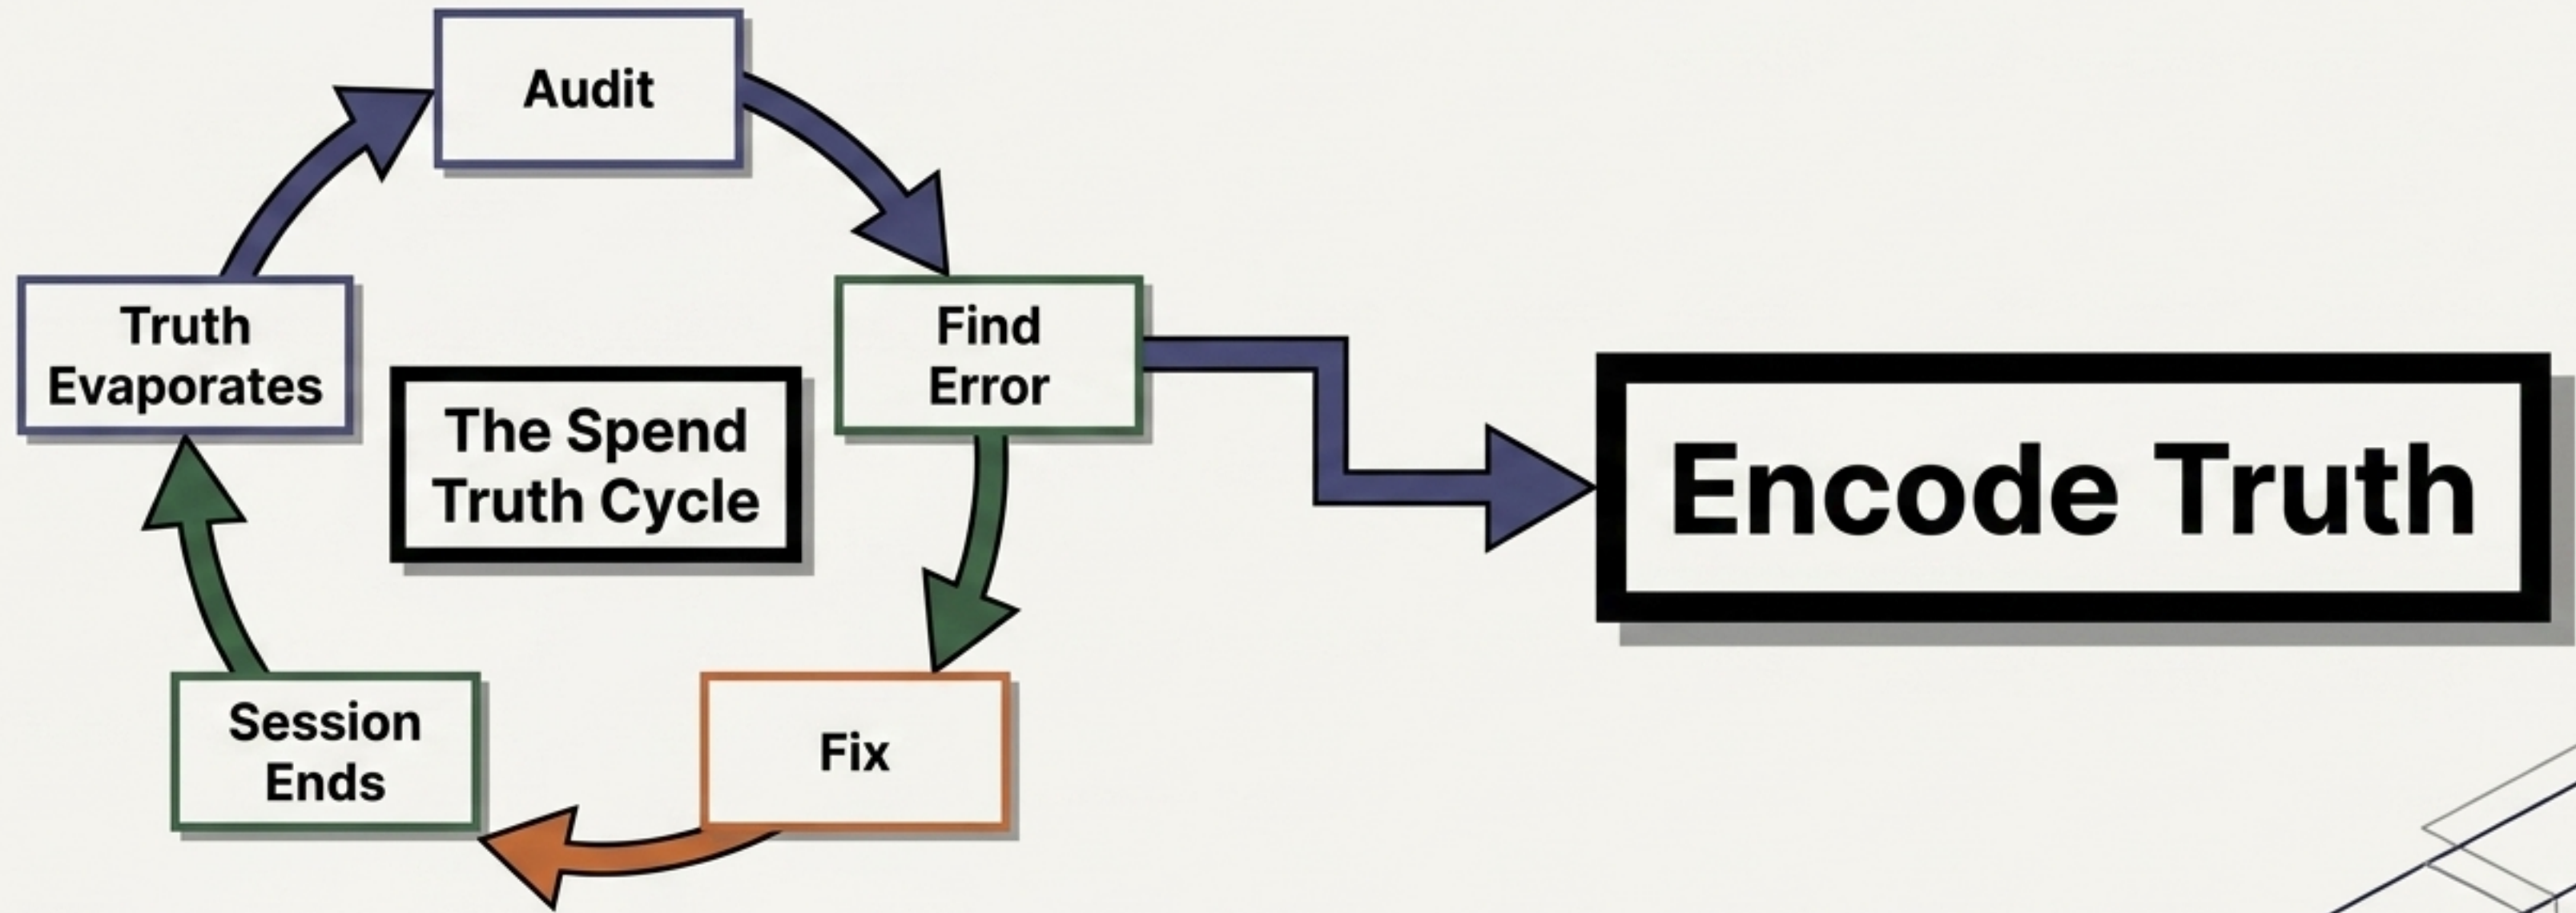
AI vibes

versioned infrastructure

Why agents reason from vibes



The verification treadmill



Every verification ends with a hard-won truth (e.g., the RFC requires wildcard support). The defensive practitioner spends that truth once, and it evaporates. Next month, a different agent makes the same error.

From spending truth to encoding it

	The Defensive Practitioner	The Compound Developer
Action	Checking claims	Grounding claims
AI Output	Produces Adjectives	Produces Verbatim Quotes
Persistence	Evaporates at session end	Versions in a corpus
Corpus Type	Latent memory	Addressable Text



The compounding move is to **encode truth** into a **versioned corpus** that every future agent reads before making claims.

Knowledge structured like a codebase


The Repository

git repository


44 regulations,
~1.4M tokens

Fragments

Citation-granularity chunks
(583 addressable units, e.g.,
GDPR Art. 28, CRA annex).



GDPR-Art.28.md



CRA-Annex.md



MiCA-Sec.5.md



DORA-Art.11.md

The Manifest

A YAML index for
deterministic navigation.

```
index: {  
  '/regulations/gdpr/art.28': 'GDPR-Art.28.md',  
  '/regulations/cra/annex': 'CRA-Annex.md'  
}
```

Provenance

Metadata tags showing consolidation
dates and official reference sources.

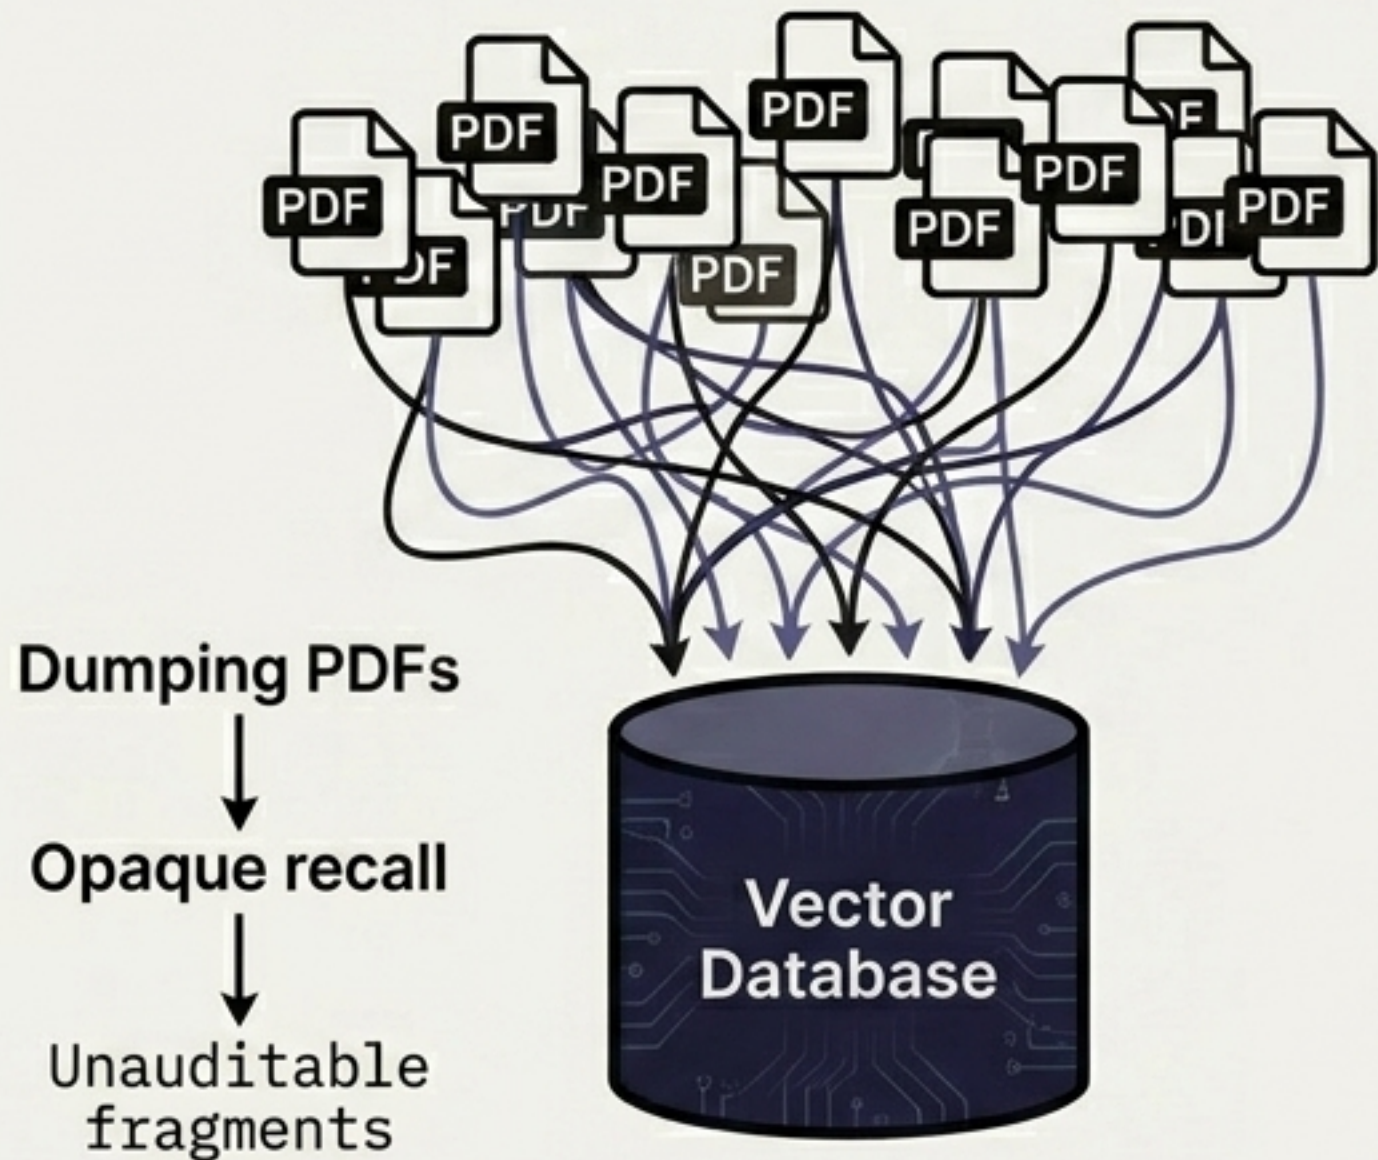
```
consolidated: '2023-10-27T14:30:00Z'
```

```
source: 'Official Journal of the European Union'
```

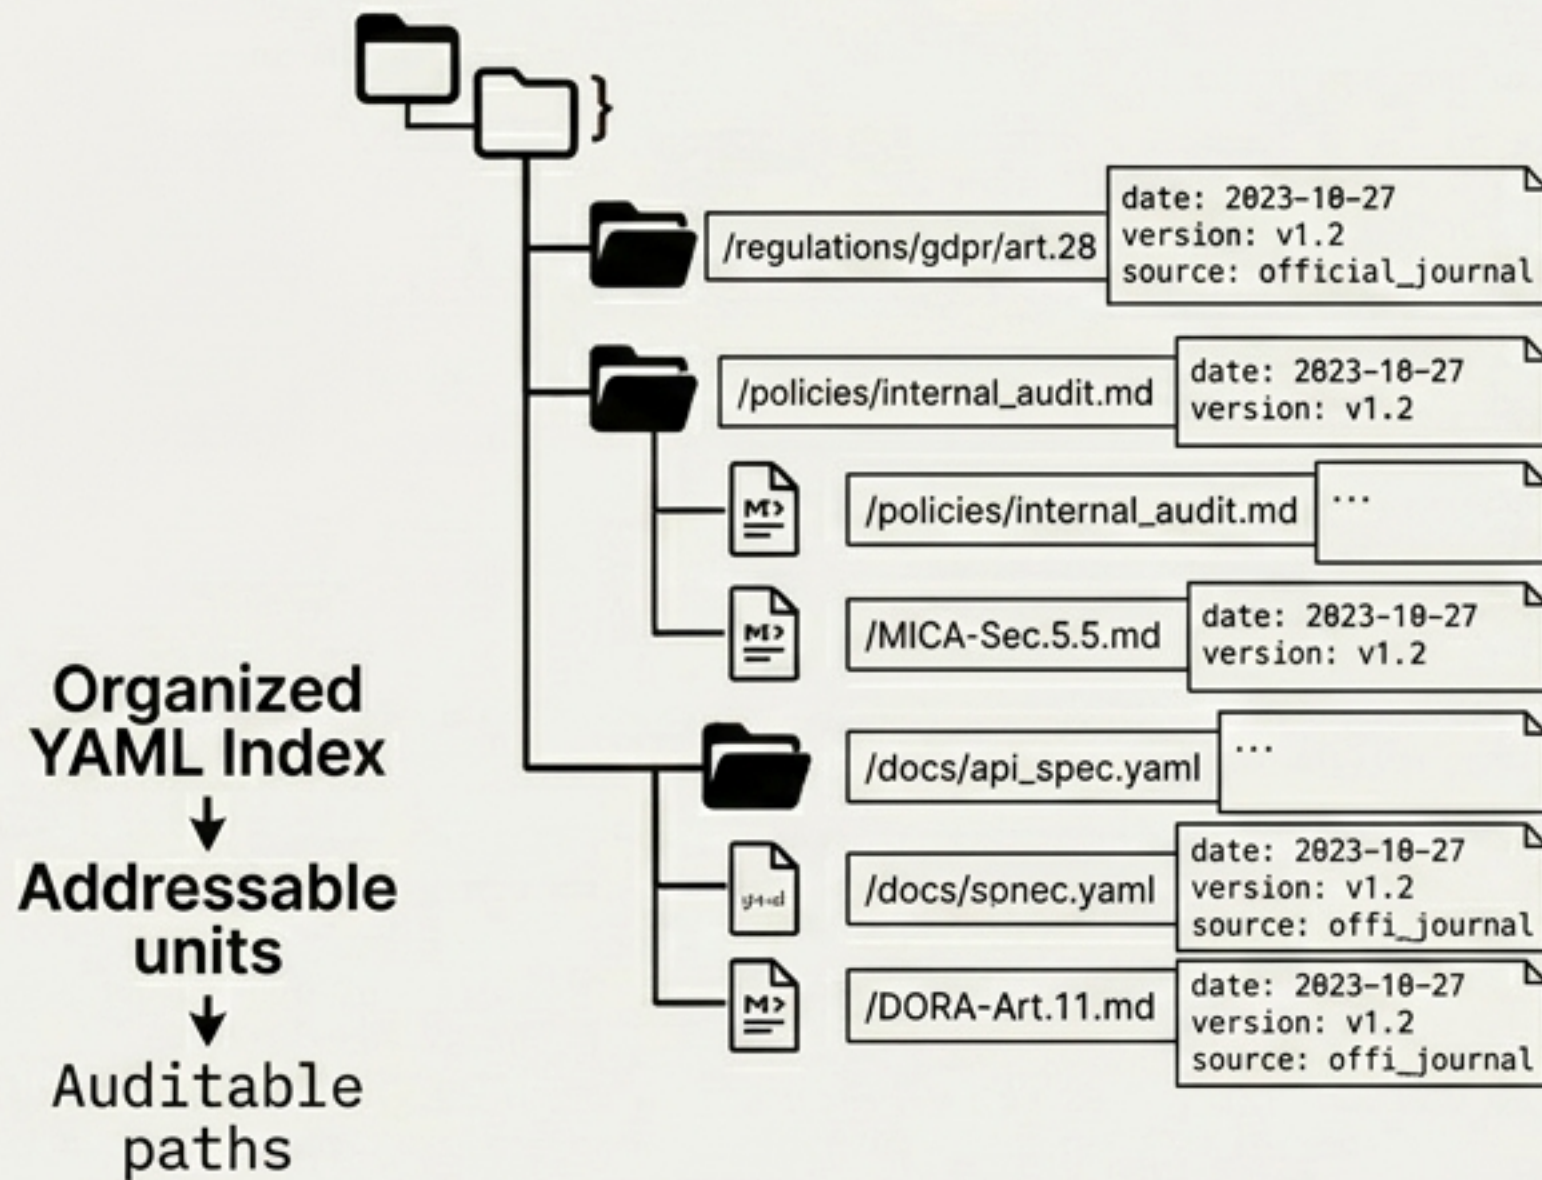
Agents navigate by manifest. Claims arrive pre-attached to the instrument you would have verified them against anyway.

Standard RAG is just vibes with extra steps

Standard RAG

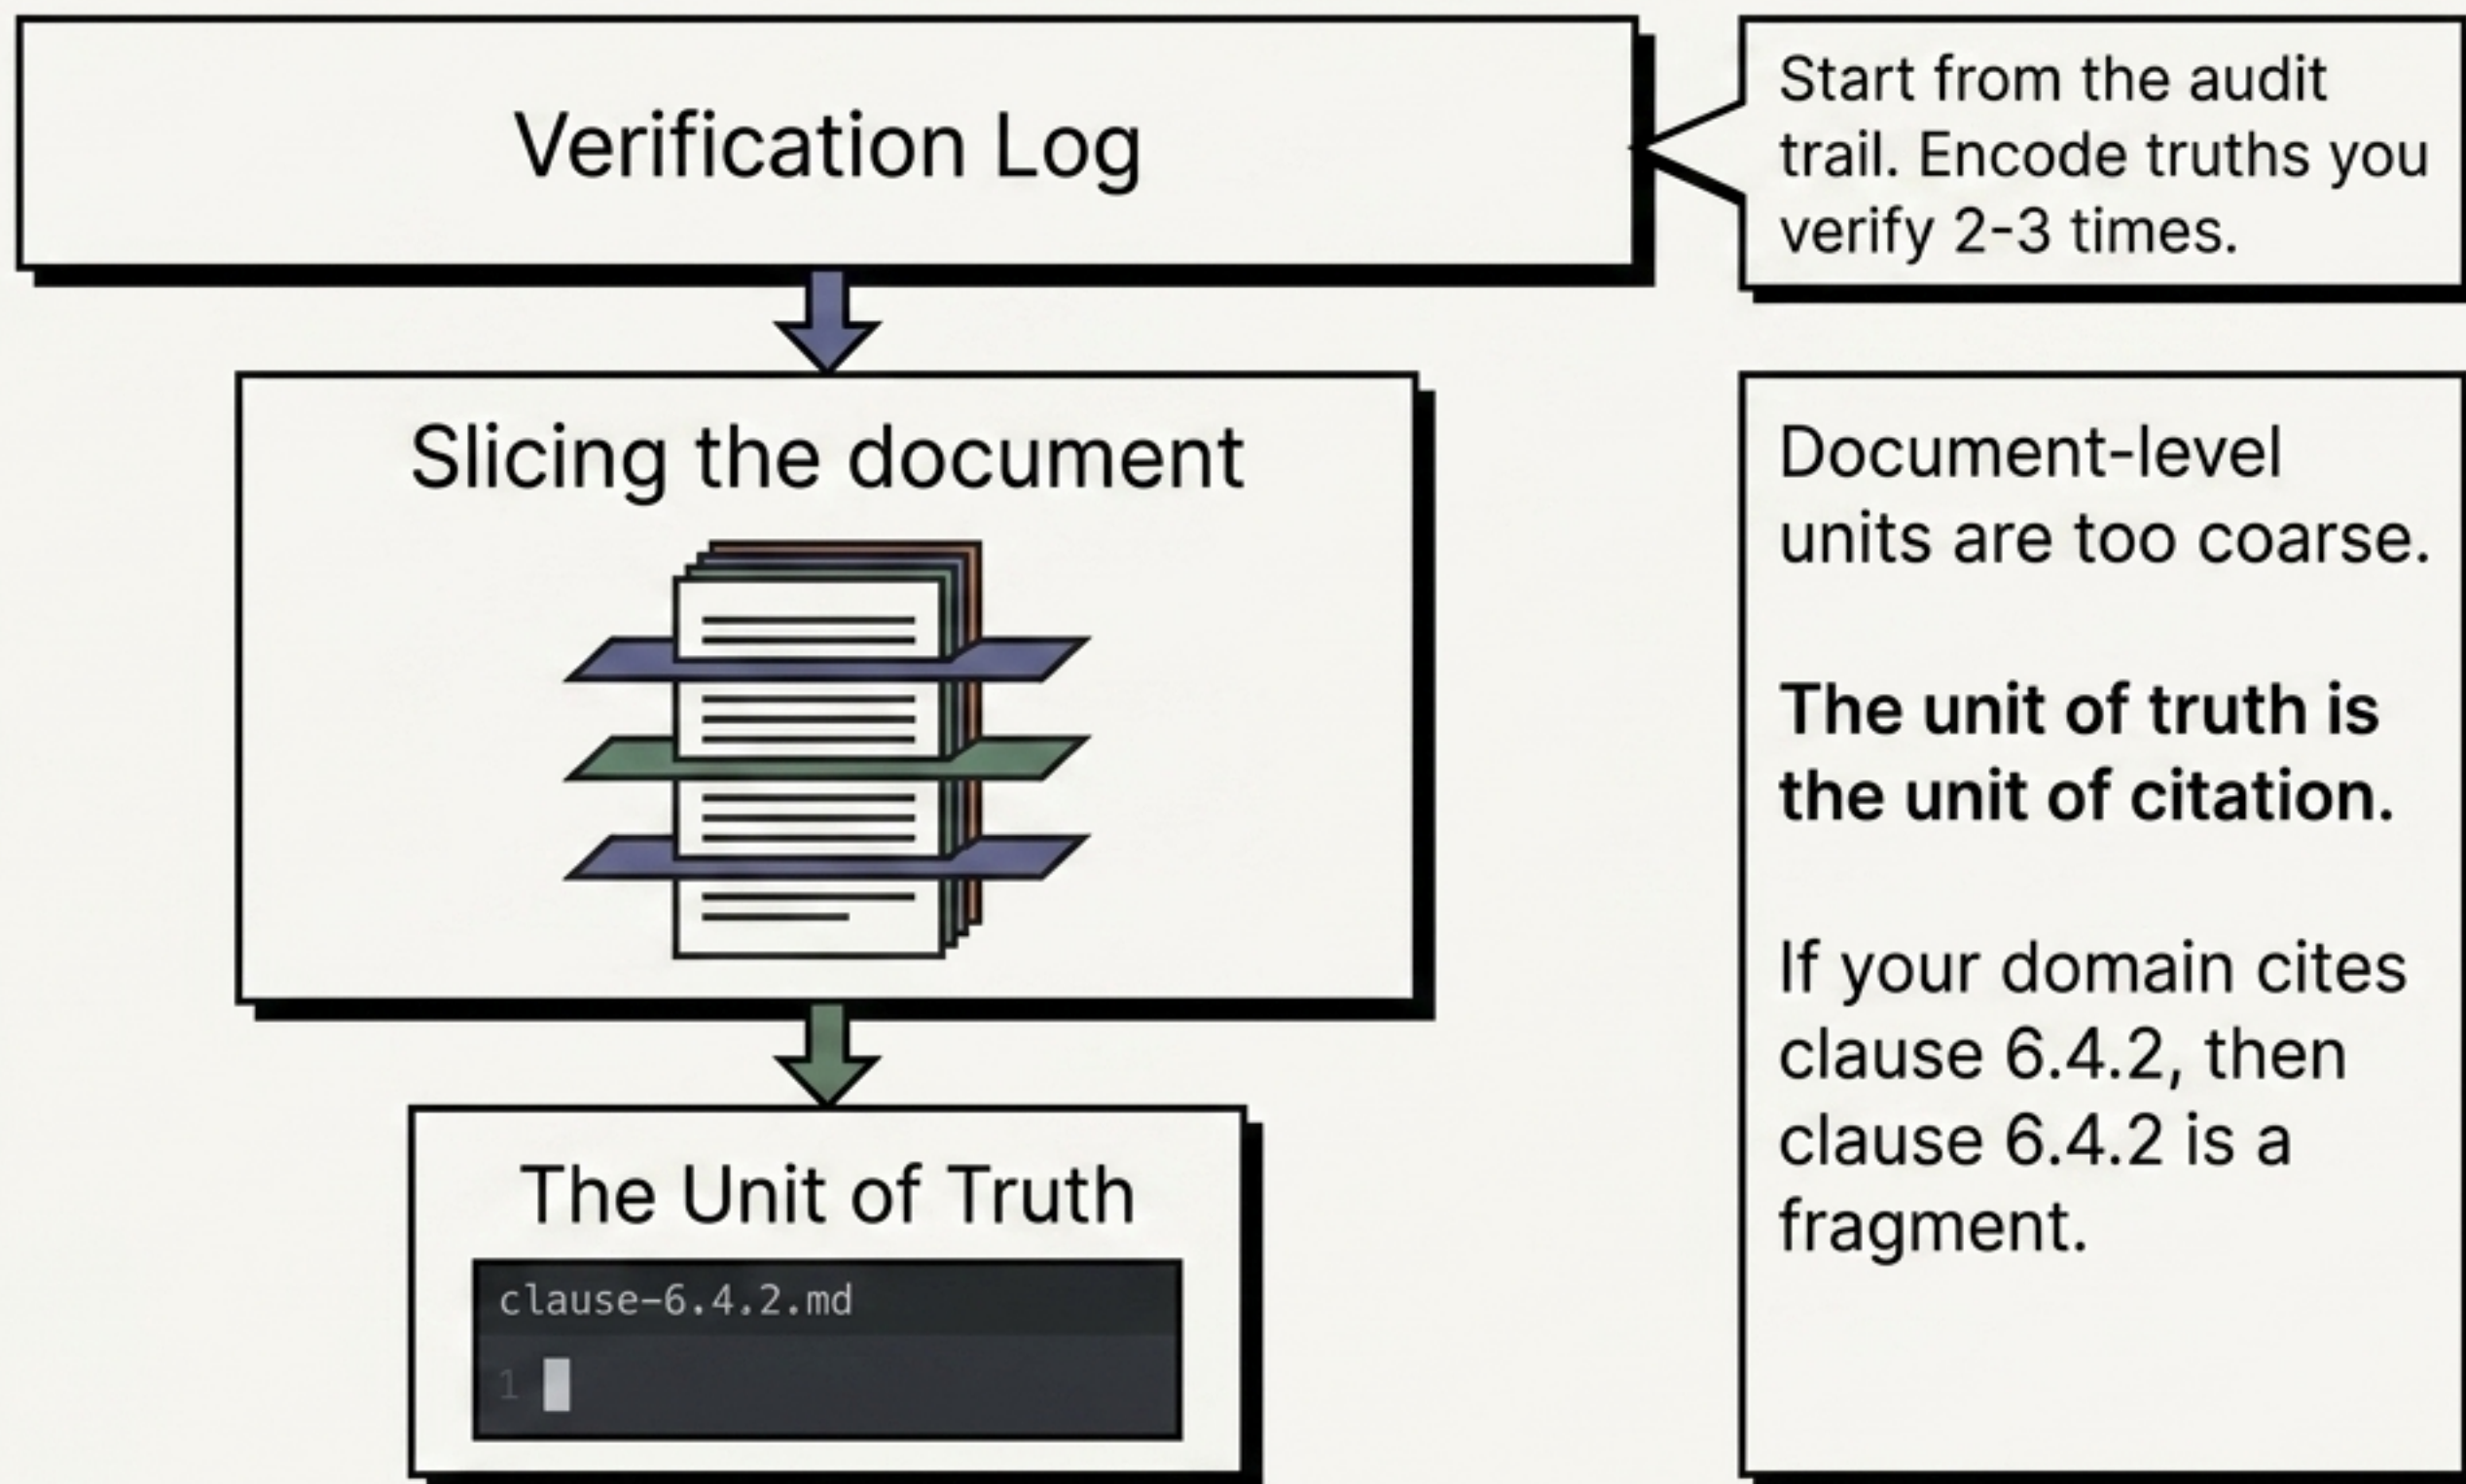


Manifest Navigation



Dumping documents into an embedding store gives you opaque recall. A manifest can be reviewed, diffed, and audited. An embedding cannot.

Targeting and fragmentation



Separating truth from interpretation

gdpr-art-28.md

v2026.04.Consolidated

[TRUTH - Immutable]

- + 120 (1) Where a controller engages a processor for carrying out
- + 121 specific or general written authorisation of the controller
- and effects or iopail suthout the legal another processor
- and nos no one, nor or not catticated object in
- authorisation of the controller.
- + 122 (2) The processor shall not engage another processor without
- + 123 specific or general written authorisation of the controller
- and effects or inpail suthout the legal another procesor
- processor shall a need to be tne permots emnloctamnon
- processors and; nor or not catticated object in
- + 123 authorisation of the controller.

internal-dpa-policy.md


Author: Admin | Date: 2026-06-11

[INTERPRETATION - Revisable]

- ← 15 This means our template needs a sub-processor clause
- 16 () Where a controller engages a processor for carrying out
- specific or general written authorisation of the controller
- and effects or iegail suthout the legal another processor
- and nos no one, nor or not caltiicated object in
- authorisation of the controller.
- 17 The processor shall not engage another processor without
- + specific or general written authorisation of the controller
- + and effects or legail suthout the legal another procesor
- + processor shall a need to be tne permose umnloctameon
- + processors and, nor or not catticated object in
- + authorisation of the controller.

Truths are immutable; interpretations encode revisable judgment.
A stale truth presented as current is a false-assurance generator.
Pin every unit to its consolidation date.

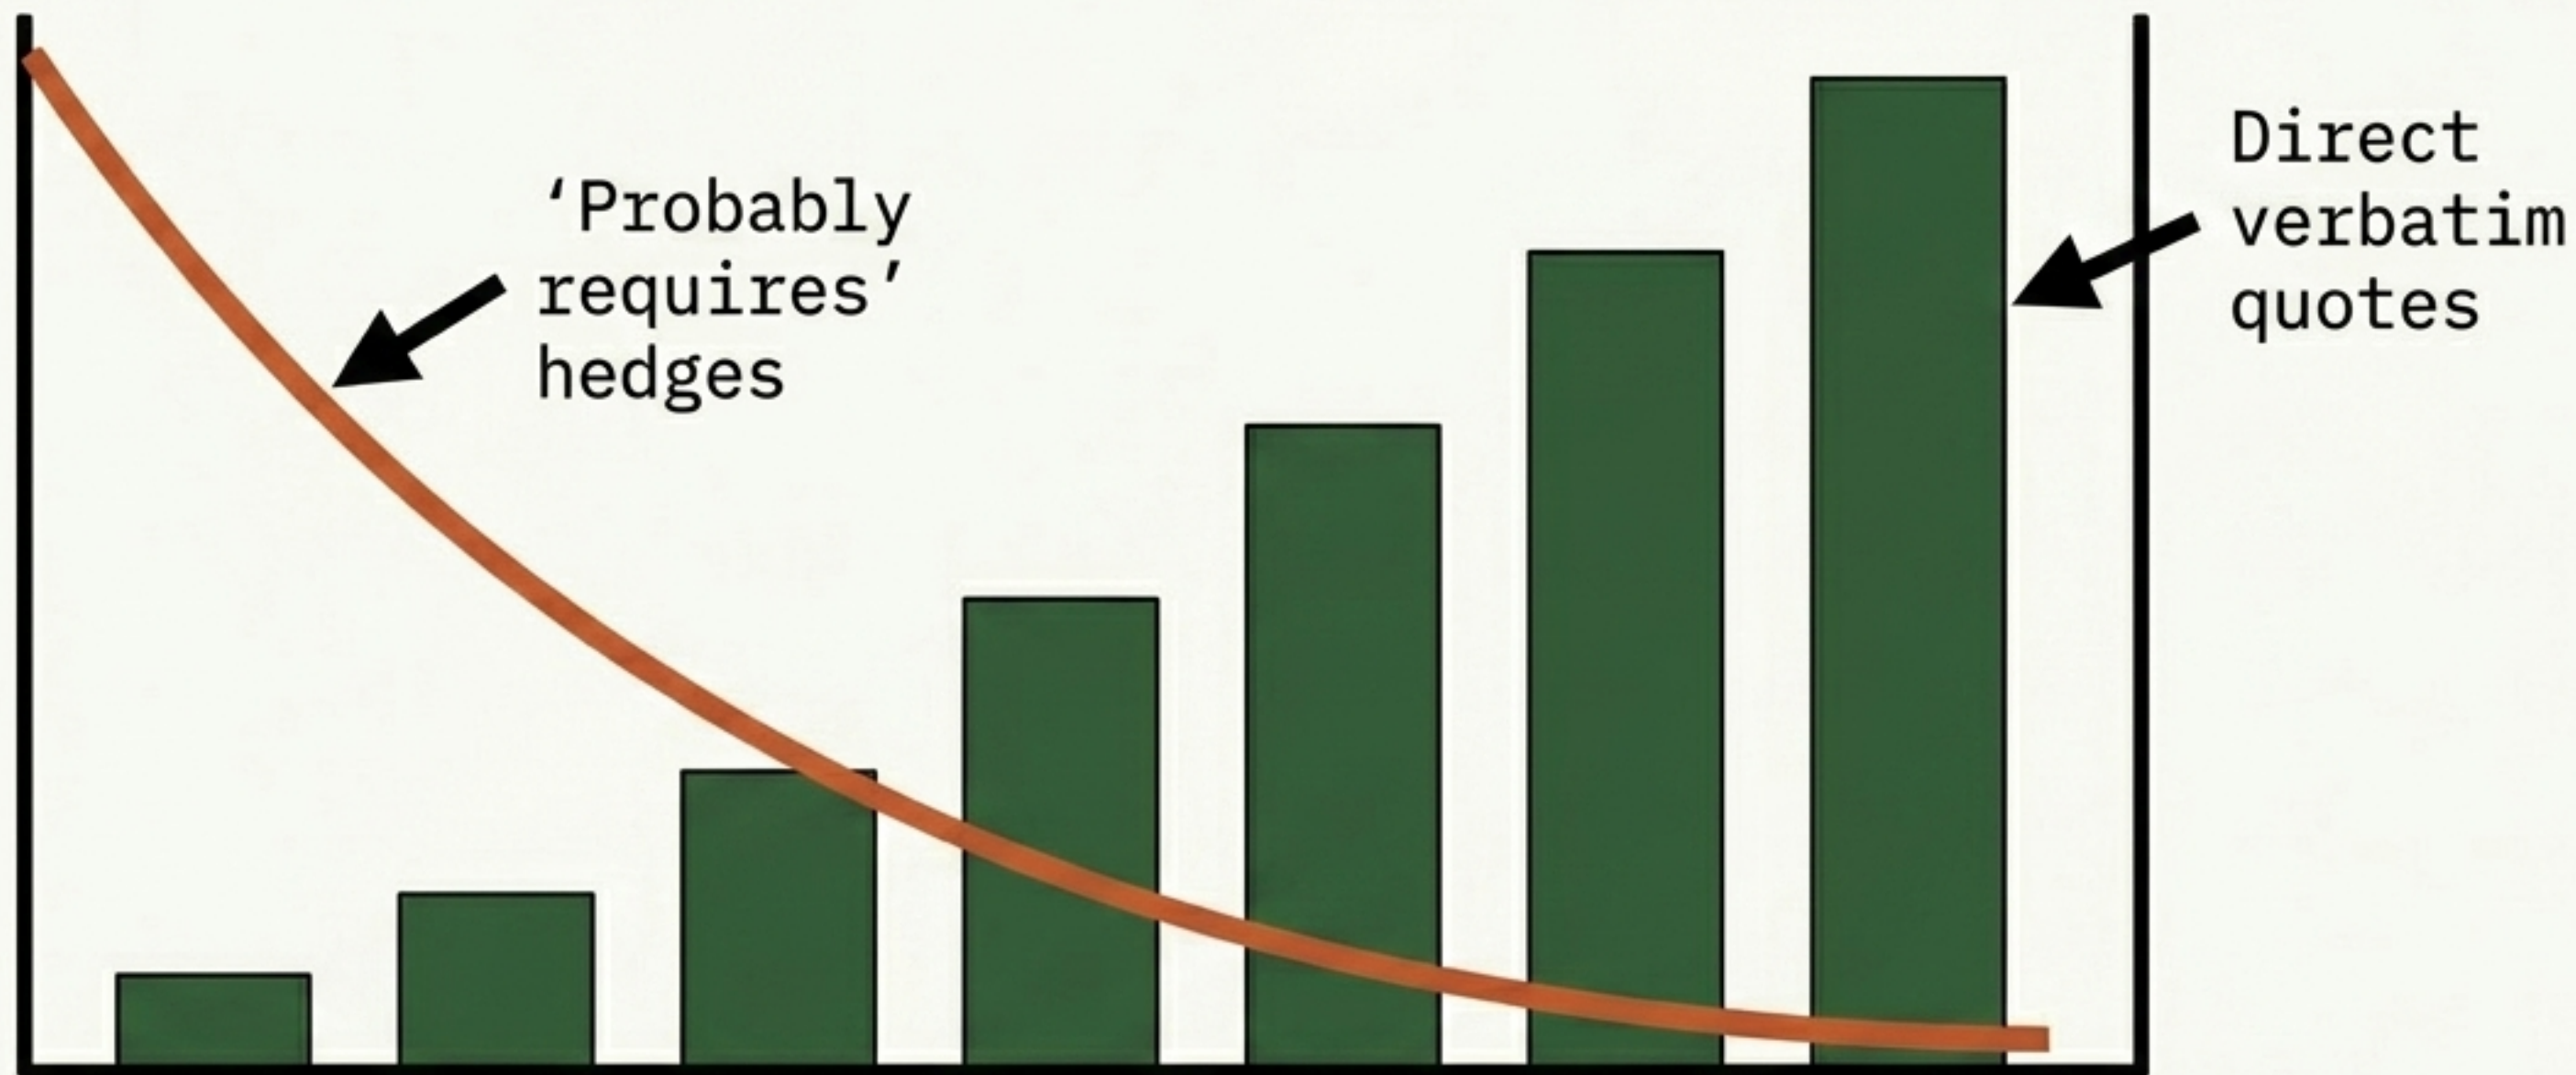
Addressable systems over searchable text



Prefer inspectability over recall. When an agent cites a `gdpr/article-28`, you can check it. When it cites 'retrieved context', you can check nothing.

The exact same principle applies to your internal system repos.

Measuring success by what stops coming back



The metric is not token counts or corpus coverage. The metric is the verification you no longer perform because categories of claims simply disappear from your audit pile.

Grounding fixes the premises, not the reasoning

[True Premise (Art. 28)] +
[Flawed Agent Logic] = [Failed Claim]

Fails in a much better-lit place.

An agent can quote an article **perfectly** and still misjudge whether a clause satisfies it. **Maintenance** is real work. But applying judgment to permanently upgrade the corpus is the real work of the **compound developer**.

The arc of the compound developer

Organized truths mean there is less to audit.

