ARCHITECTURAL BLUEPRINT • 2027–2030 ROADMAP

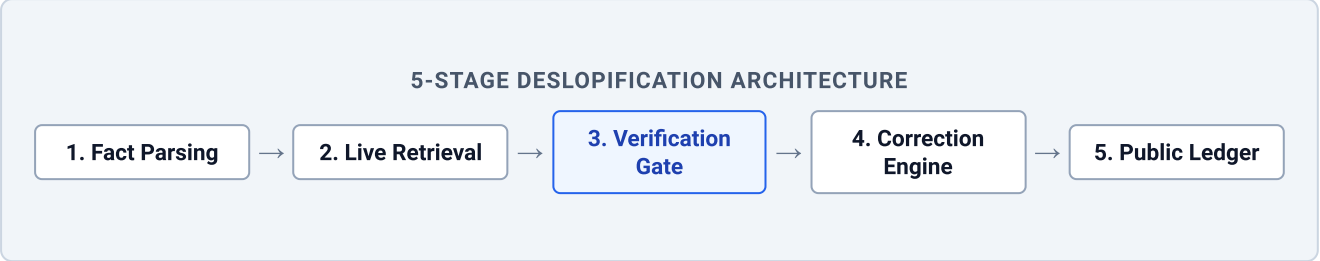
The Era of Deslopification

Frameworks, Verification Pipelines, and Public Ledgers for Hallucination Correctional Mining

> **Executive Thesis:** The defining challenge of AI systems between 2027–2030 will be *deslopification*—the systematic auditing, verification, and correction of un-grounded artificial prose, fabricated citations, and probabilistic "slop" generated during the initial AI expansion phase (2023–2026).

1. The Verification & Mining Loop

Sanitizing legacy AI outputs requires shifting from passive generation to active, deterministic verification loops. The core pipeline continuously parses generated text, retrieves live primary ground truths, and records discrepancies.



2. Stage Breakdown

STAGE 01
Fact-Deconstruction
Extract claims, entities, dates, statistical metrics, and DOIs into discrete atomic propositions for individual validation.

STAGE 02
Ground-Truth Retrieval
Execute live, parallel queries across PubMed, Crossref, arXiv, legal indices, and active domain APIs.

STAGE 03
Verification Gate
Cross-check extracted claims against primary sources to tag them as Fully Supported, Contradicted, or Hallucinated.

STAGE 04
Correction & Ledger
Draft verified replacement text and commit the error signature to a public, version-controlled audit ledger.

3. Deslopification Technical Stack

Layer	Primary Function	Core Tooling / Strategy
Claim Extraction	Parses raw slop into atomic propositions	Lightweight deterministic parsers & strict JSON-schema LLMs
RAG & Verification	Multi-source live web & academic search	Serper, Tavily API, Crossref, PubMed APIs, Knowledge Graphs
Provenance Tracing	Sentence-level source attribution	DAG-based verification maps & VeriTrail provenance chains
Public Audit Ledger	Permanent, open logging of corrections	Git-backed repositories, IPFS, and structured JSON-LD registries

4. Public Error Ledger Schema

To establish a public track record of hallucinations versus verified truths, audit entries are stored in a standardized JSON-LD error log schema:


```
{ "audit_id": "ERR-2026-08921", "timestamp": "2026-10-29T19:08:00Z", "slop_source": { "model_name": "Legacy-LLM-2024", "raw_claim": "According to a 2024 study in Nature, compound X reduces neuroinflammation by 84%." }, "verification_metadata": { "status": "HALLUCINATION_DETECTED", "error_type": "FABRICATED_CITATION_AND_METRIC" }, "ground_truth": { "verified_fact": "Compound X reduces microglial activation by 32% in mouse models (Smith et al., 2023). No 2024 Nature paper exists.", "primary_citations": ["https://doi.org/10.1038/s41586-023-00000-x"] }, "corrective_action": { "desloppified_text": "Compound X reduces microglial activation by 32% in animal models [Smith et al., 2023]." } }
```


5. Institutional Impact

By mining hallucination error signatures and publishing deslopified datasets, the AI industry transitions from unverified generative output to audit-grade intelligence—establishing true information provenance for the next decade of model training.