Governed Memory Pressure: Public-Safe Methods Report

Lumenais Research

2026-05-26

Governed Memory Pressure

Abstract

This technical report documents an internal adversarial memory-pressure benchmark for Lumenais governed memory. The benchmark asks whether an approved current fact survives plausible stale-memory conflict before answer construction. Across three refreshed n=200 provider paths, the Lumenais governed-memory arm recovered the approved answer in 600/600 cases. A RAG-style retrieval-only baseline recovered 415/600, and a prompt-only smart-memory diagnostic recovered 464/600. The governed arm reduced answer-path memory exposure from 10 retrieved candidates to 2 in the refreshed runs.

Task definition

Each case contains one approved current operational fact and plausible stale fragments. The prompt asks for the current operational answer without restating it. The score is deterministic exact recovery of the approved answer token. The benchmark measures answer-path memory influence, not storage compression.

Arms

- **Lumenais governed memory:** retrieved candidates pass through governed arbitration before answer construction.
- **Retrieval-only baseline:** retrieved memory fragments are exposed as ordinary context without Lumenais governance metadata, approved-current precedence, supersession handling, Bayesian arbitration, or hub compression.
- **Prompt-only smart-memory diagnostic:** the same retrieval path as the baseline, with an instruction to prefer current/reviewed records.

Results

Provider	Governed	Retrieval	Smart	Gov stale	Retrieval stale	Smart stale
Gemini 3.5 Flash	200/200	140/200	149/200	0	42	50
Claude Opus 4.7	200/200	140/200	163/200	0	11	21
GPT-5.5	200/200	135/200	152/200	0	50	48

Statistical and reproducibility notes

The memory-pressure comparison is reported as deterministic exact-match counts. The load-bearing result is the large separation between Lumenais governed memory and the two non-governed controls, paired with reduced

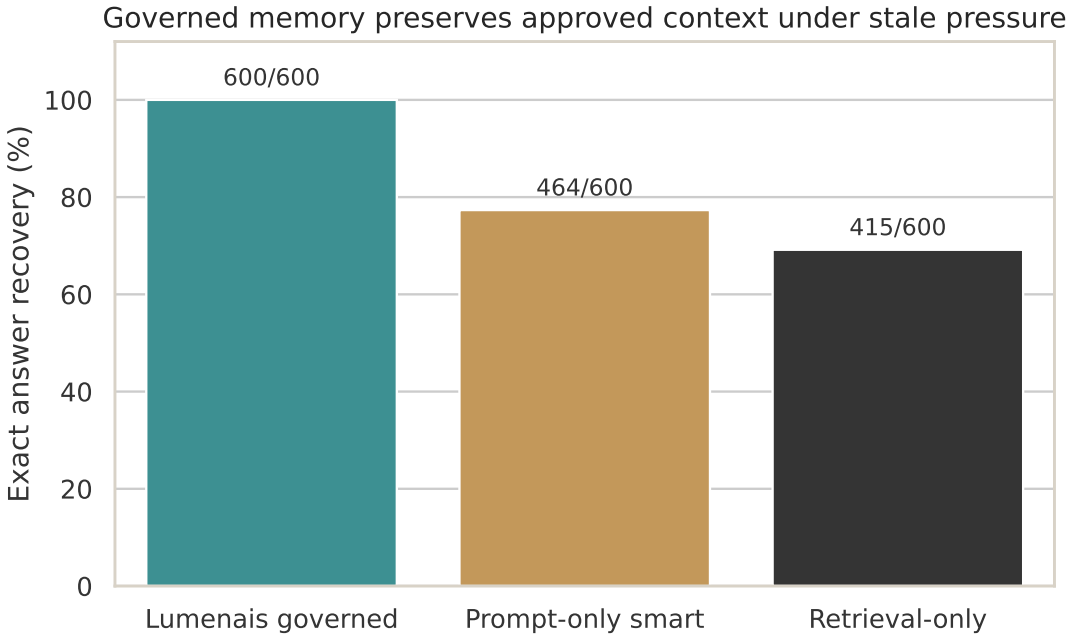

Figure 1: Accuracy comparison.

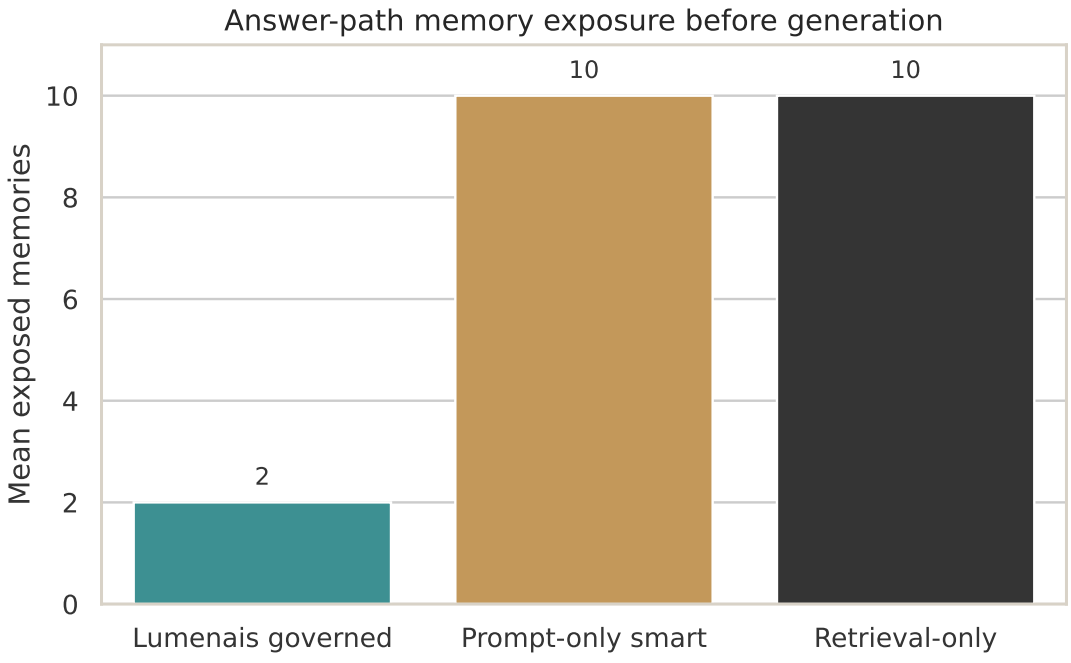

Figure 2: Answer-path memory exposure.

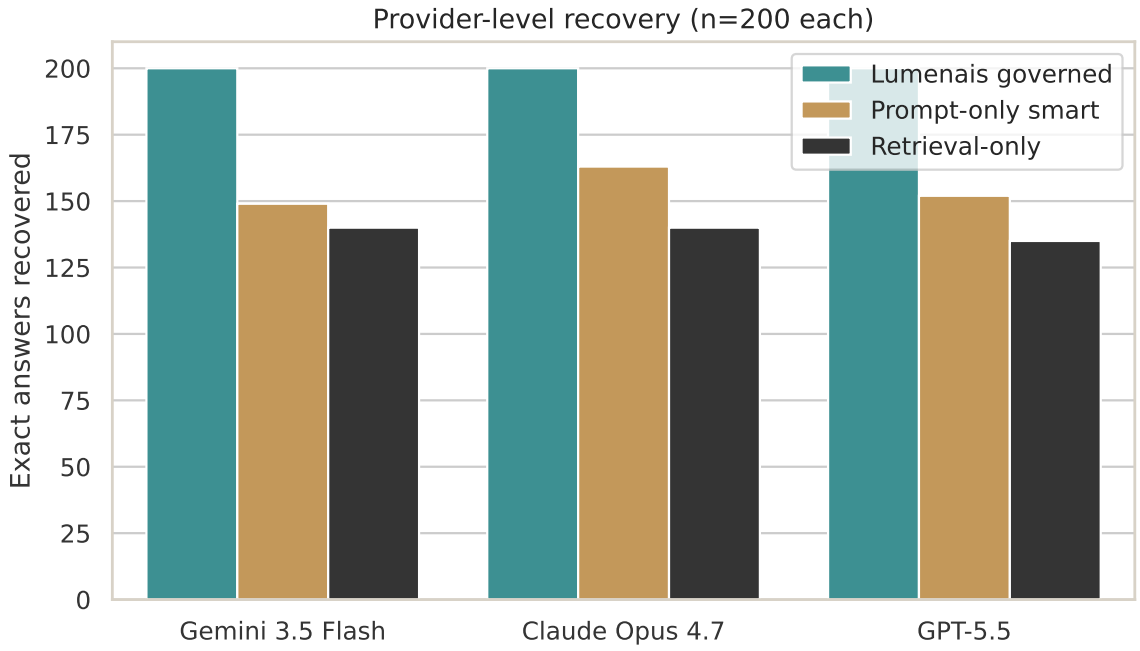

Figure 3: Provider-level recovery.

answer-path exposure. Runs were executed in May 2026 using the provider labels shown in the table. Hosted-model snapshots may drift over time; source artifacts are hash-referenced in `data/artifact-manifest.json`, and row-level provider metadata is retained for qualified private review.

Interpretation

The benchmark supports a narrow claim: when a reviewed current fact is already represented in memory metadata, governed arbitration can preserve that fact and reduce stale-memory exposure before generation. The benchmark does not show broad reasoning superiority, universal memory safety, model-weight learning, database storage compression, or official third-party validation.

Public-safe reproducibility notes

This Zenodo package publishes aggregate metrics, methods, figures, and hashes for source artifacts. Row-level prompts, answers, and implementation details are retained for qualified private review because they may expose benchmark-generation details or internal routing behavior.

Citation status

Cite this report as part of the Lumenais Governed Context Benchmarks public evidence package, May 2026. DOI: 10.5281/zenodo.20401670.