

Corbenic AI Releases Technology That Eliminates AI's Largest Cost

Save AI memory on one server, restore it byte-identical on another, even across GPU generations. Cryptographic proof. Long-context AI up to 21× faster.

DE, UNITED STATES, June 4, 2026 /EINPresswire.com/ -- Corbenic AI Releases Memory Engine That Eliminates Redundant Recomputation



You don't need a bigger brain. You need a better memory."

Schelpe Sietse

Corbenic AI today announced Taliesin, a memory engine that eliminates the single largest recurring cost in enterprise AI deployments: redundant recomputation of context a model has already processed.

Every time an AI system answers a question about a

document, it re-reads that document from scratch. Across ten queries on a 100-page report, the model processes a thousand pages. Taliesin saves what the model has already read and restores it on demand — mathematically identical to a fresh re-read, down to the last bit.

Two minutes becomes seven seconds.

On a \$0.69-per-hour graphics card, the longest test contexts took a model more than two minutes to process from scratch. Taliesin restored them in under seven seconds: a 21-times speedup, with no loss of accuracy.

Corbenic validated the approach across GPU generations. In a bidirectional relay between an Ampere A6000 and an Ada Lovelace RTX 4090, Taliesin moved AI memory back and forth and produced 64 of 64 output tokens identical to what the originating card would have generated. Cross-architecture verification of this kind is seldom attempted publicly: GPU generations round floating-point operations differently, and most disaggregated-serving systems accept approximation as sufficient. A context prefilled on a lower-cost card and served from a higher-performance one produces the same answer either way.

Cryptographic verification, public and reproducible.

Corbenic published SHA-256 hashes — the same fingerprint standard used to verify software downloads — for every trial. Any researcher can run the verification suite on three public open-weight models from Meta, Alibaba, and Mistral and confirm the hashes match. Results: 45 of 45 trials matched. A 60-process follow-up: 60 of 60. Cross-machine between two physically separate servers: matching in both directions.

A EUR 600 model, by design.

Alongside Taliesin, Corbenic released Galahad-0.5B, an open-source 570-million-parameter model trained from scratch for EUR 600 (\$677) — roughly the cost of a mid-range laptop. Galahad does not outperform similar-sized open models on standard benchmarks; its purpose is to make the full verification chain auditable, weight by weight.

A layered cost-reduction stack.

Taliesin pairs with Merlin, Corbenic's open-source byte-exact deduplication engine (available on arXiv and GitHub). Merlin removes redundant tokens before compute; Taliesin eliminates recompute on every subsequent reuse. Together, the two components attack recurring AI compute from both ends: Merlin delivers 13.9 to 71 percent input reduction; Taliesin delivers up to 21 times less compute per reused context. For reuse-heavy workloads, the combined reduction compounds to well over 90 percent of the recurring compute bill.

"The industry has focused on building bigger models. Corbenic focused on building better memory," said Sietse Schelpe, Founder and CEO of Corbenic AI. "Taliesin emerged from validation work on Merlin — in the process of proving deduplication worked end to end, the engine for eliminating recomputation became apparent."

Availability

Galahad-0.5B (open-source): huggingface.co/Corbenic/Galahad-0.5B-base

Merlin Community Edition: github.com/corbenicai/merlin-community

Verification receipts: huggingface.co/datasets/Corbenic/taliesin-receipts

Technical paper / design-partner program: sietse@corbenic.ai

About Corbenic AI

Corbenic AI is a Delaware-incorporated company headquartered in Oostkamp, Belgium, building infrastructure that reduces the cost of running AI without reducing its accuracy. More at corbenic.ai.

Sietse Schelpe

Corbenic AI

[email us here](#)

Visit us on social media:

[LinkedIn](#)

[X](#)

[Other](#)

This press release can be viewed online at: <https://www.einpresswire.com/article/917089794>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.