

# One AI architecture produced a structured decision-support package for a family navigating dementia in 1 hour 33 minutes — and told the user a hard truth they did not want to hear.

Thinking Studios' operating system, on the kind of life-altitude question current AI systems usually deflect.

## 1 h 33m 38s

From a single one-sentence prompt to a structured decision-support package for a family in early-stage dementia diagnosis: a four-option comparison with emergent fourth option, a Month-1 critical-path callout, a financial trajectory across 10 years, a clinical contradiction worth a doctor's attention, and three protocol designs (preference elicitation, family meeting, sibling-conflict diagnostic). Every claim provenance-linked to the audit trail that produced it.

### 190

TOOL INVOCATIONS / 10  
DISTINCT TOOLS

### 20

SPECIALIST CHILD THOUGHTS  
IN 7 WAVES

### 54

LAYER-1 INTERPRETATIONS,  
CALIBRATED

### 132

PROVENANCE LINKS  
RECORDED

### 19 / 9 / 0

CONVERGENCES /  
DIVERGENCES / UNRESOLVED

### 4

OPTIONS ANALYZED (1  
EMERGENT, NOT IN ORIGINAL  
FRAMING)

### \$1.39M vs \$2.81-3.46M

10-YEAR FINANCIAL  
TRAJECTORY (BEST VS  
PREFERRED)

### 5

SECTION REPORT (HEADING ·  
CRITICAL-PATH · MATRIX ·  
FINDINGS · ACTION LIST)

#### THE LOAD-BEARING FINDING

"The user's stated bandwidth (2-7 hours per week) cannot meet the actual demand of moderate dementia (42-100+ hours per week). Both in-home care and move-in care become caregiver-collapse trajectories on this gap. Projected failure point: months 8-14. This is not a matter of love or commitment; it is an operational arithmetic."

— derived from five independent specialist strands that converged on the same finding. Telling someone "the options you prefer are infeasible" is hard. Umma produced it because the cross-strand evidence supported it.

## WHAT UMMA IS · WHY THIS CASE MATTERS

Umma is an **operating system** — not a wrapper around a foundation model, not an autonomous-agent framework. The architecture decomposes work into coordinated waves of specialist child thoughts, reconciles their outputs through adversarial synthesis, records its own reasoning convergences and divergences as a queryable audit trail, and refuses to fabricate when grounding fails. Built by Thinking Studios, in development since early 2024. The run shown here was produced by Umma v2; Wave 1 of v3 deploys the week of **June 8, 2026**.

Current AI systems usually answer questions like "my mother was diagnosed with dementia" with either generic resource lists or with emotionally-cautious deflection. Both responses fail the family. What this family needed was **structured decision architecture** — the option matrix, the bandwidth audit, the financial model, the protocol designs, the clinical-contradiction flag. Umma did that. The architecture is what made it possible. Companion to this work: the Netflix Lemur audit, which demonstrates the same architecture on an institutional question.

#### CONTACT

Kim Jeong Ha, founder  
Thinking Studios  
kjh@thinkingstudios.co

#### VERIFY THE RUN

- Full reasoning trail and methodology walkthrough: included as PDF attachments in the press distribution package
- Reproducibility documentation: specific database queries reproduce every cited number; attached as PDF
- The verbatim assistant message Umma returned to the user; the executive summary; the 5-section delivered report — all attached