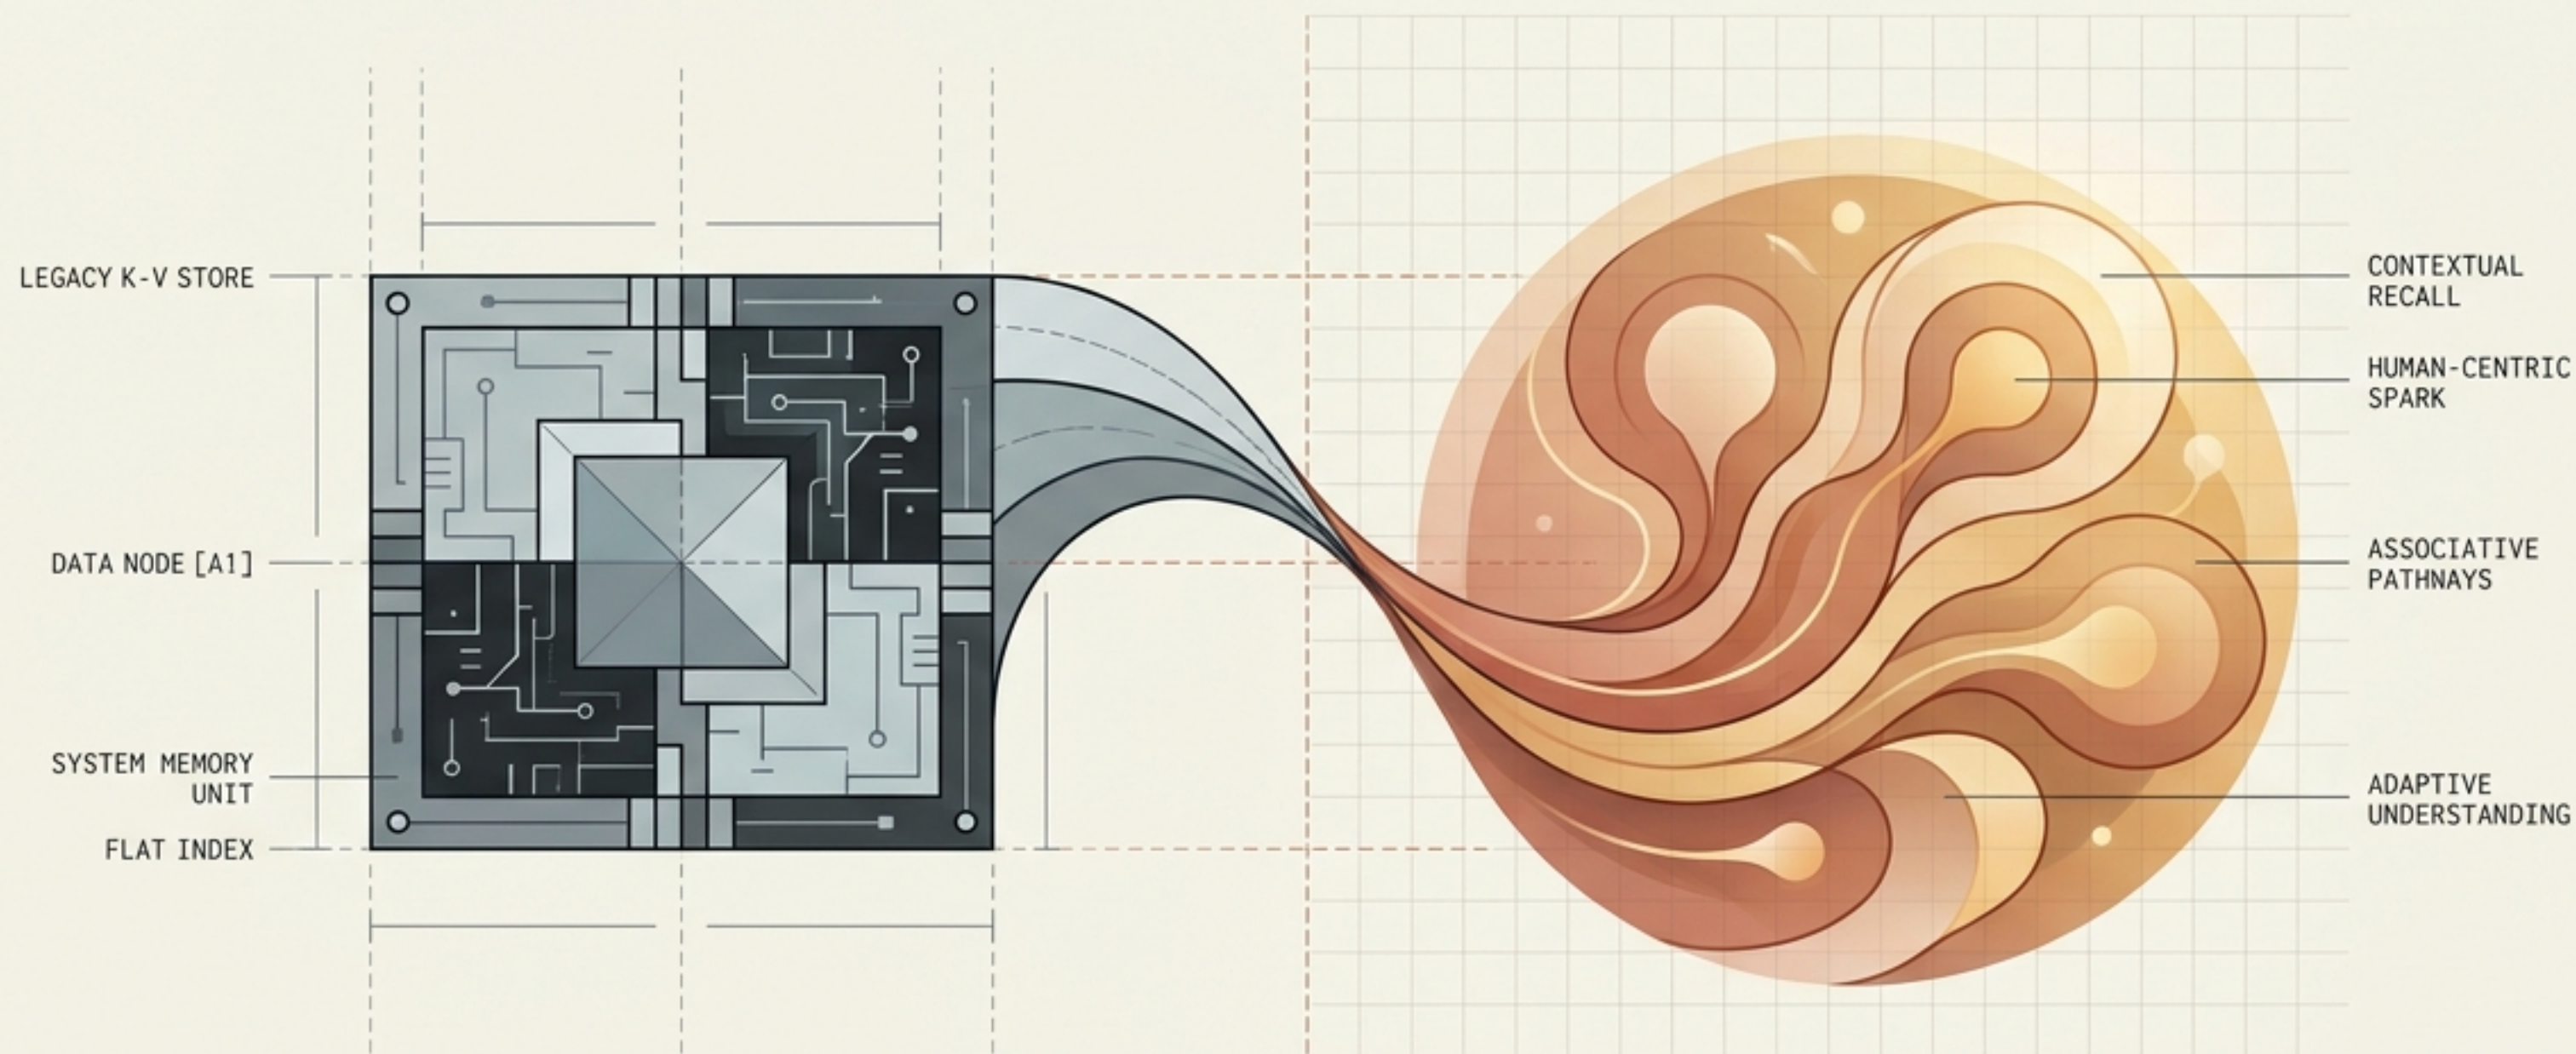
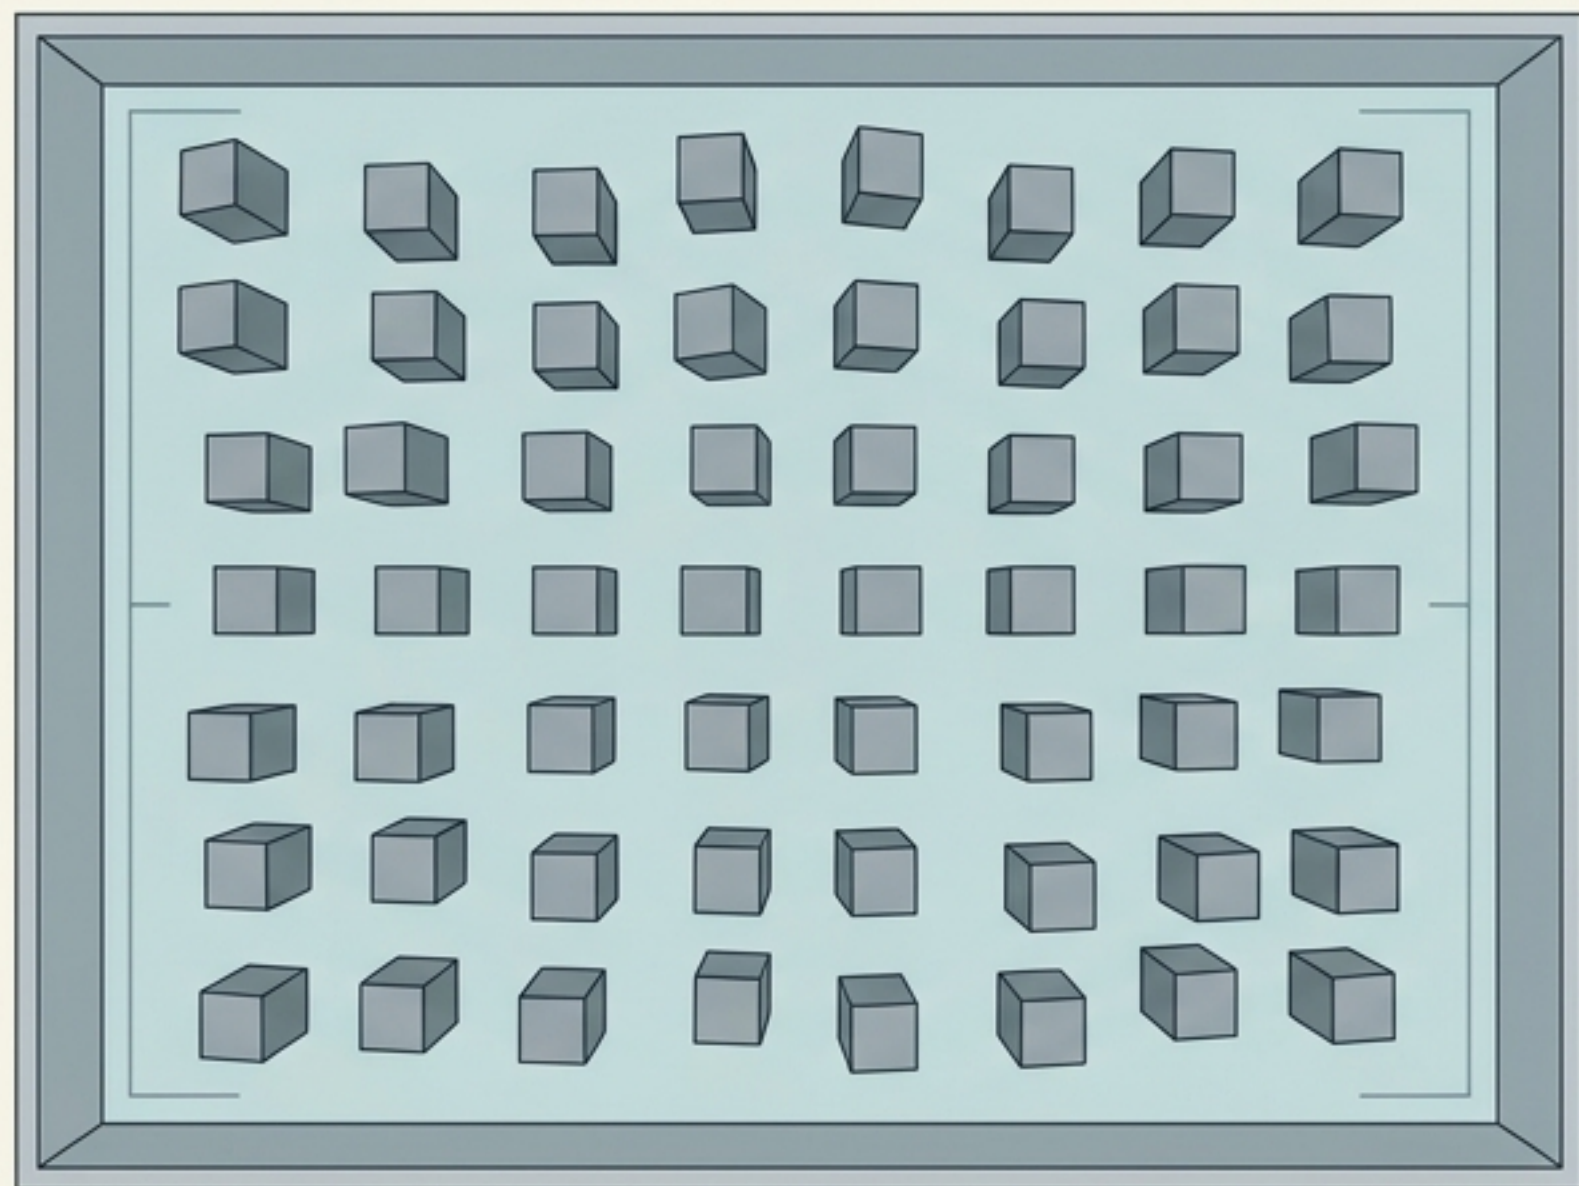


Memory That Feels Human.

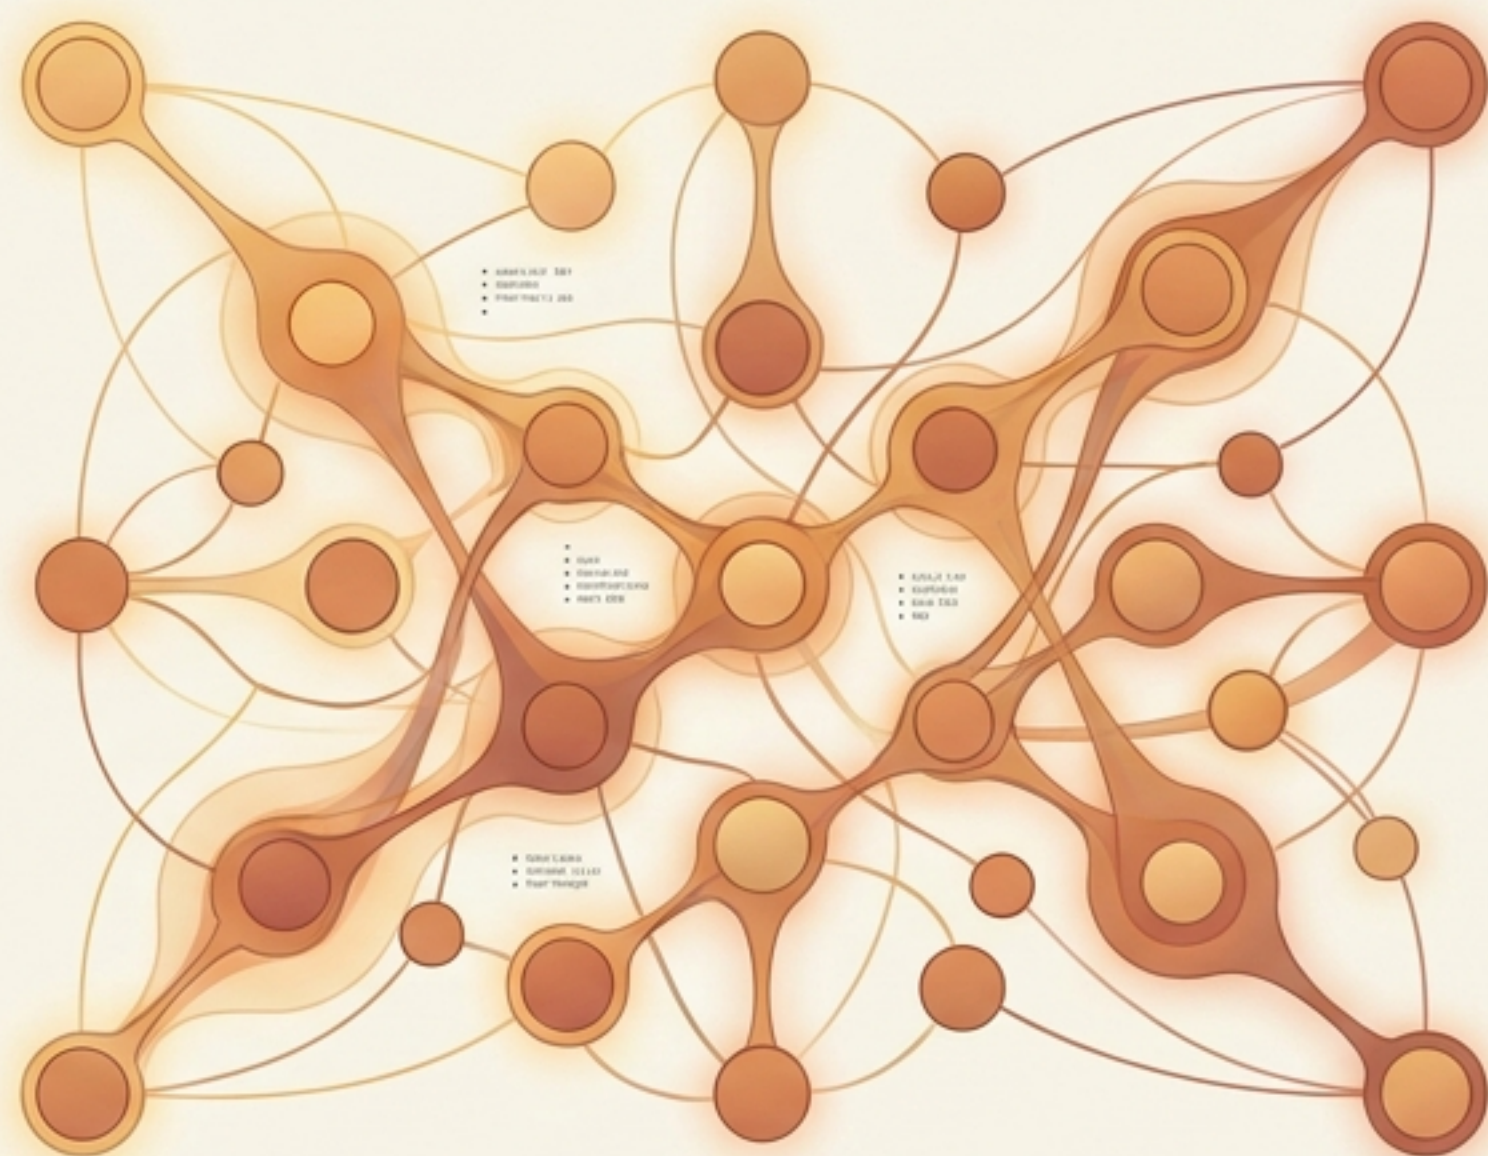
Rebuilding Mio's memory from a flat key-value store to contextual recall: A Three-Phase Architectural Blueprint.



Chatbot with Amnesia




True Companion



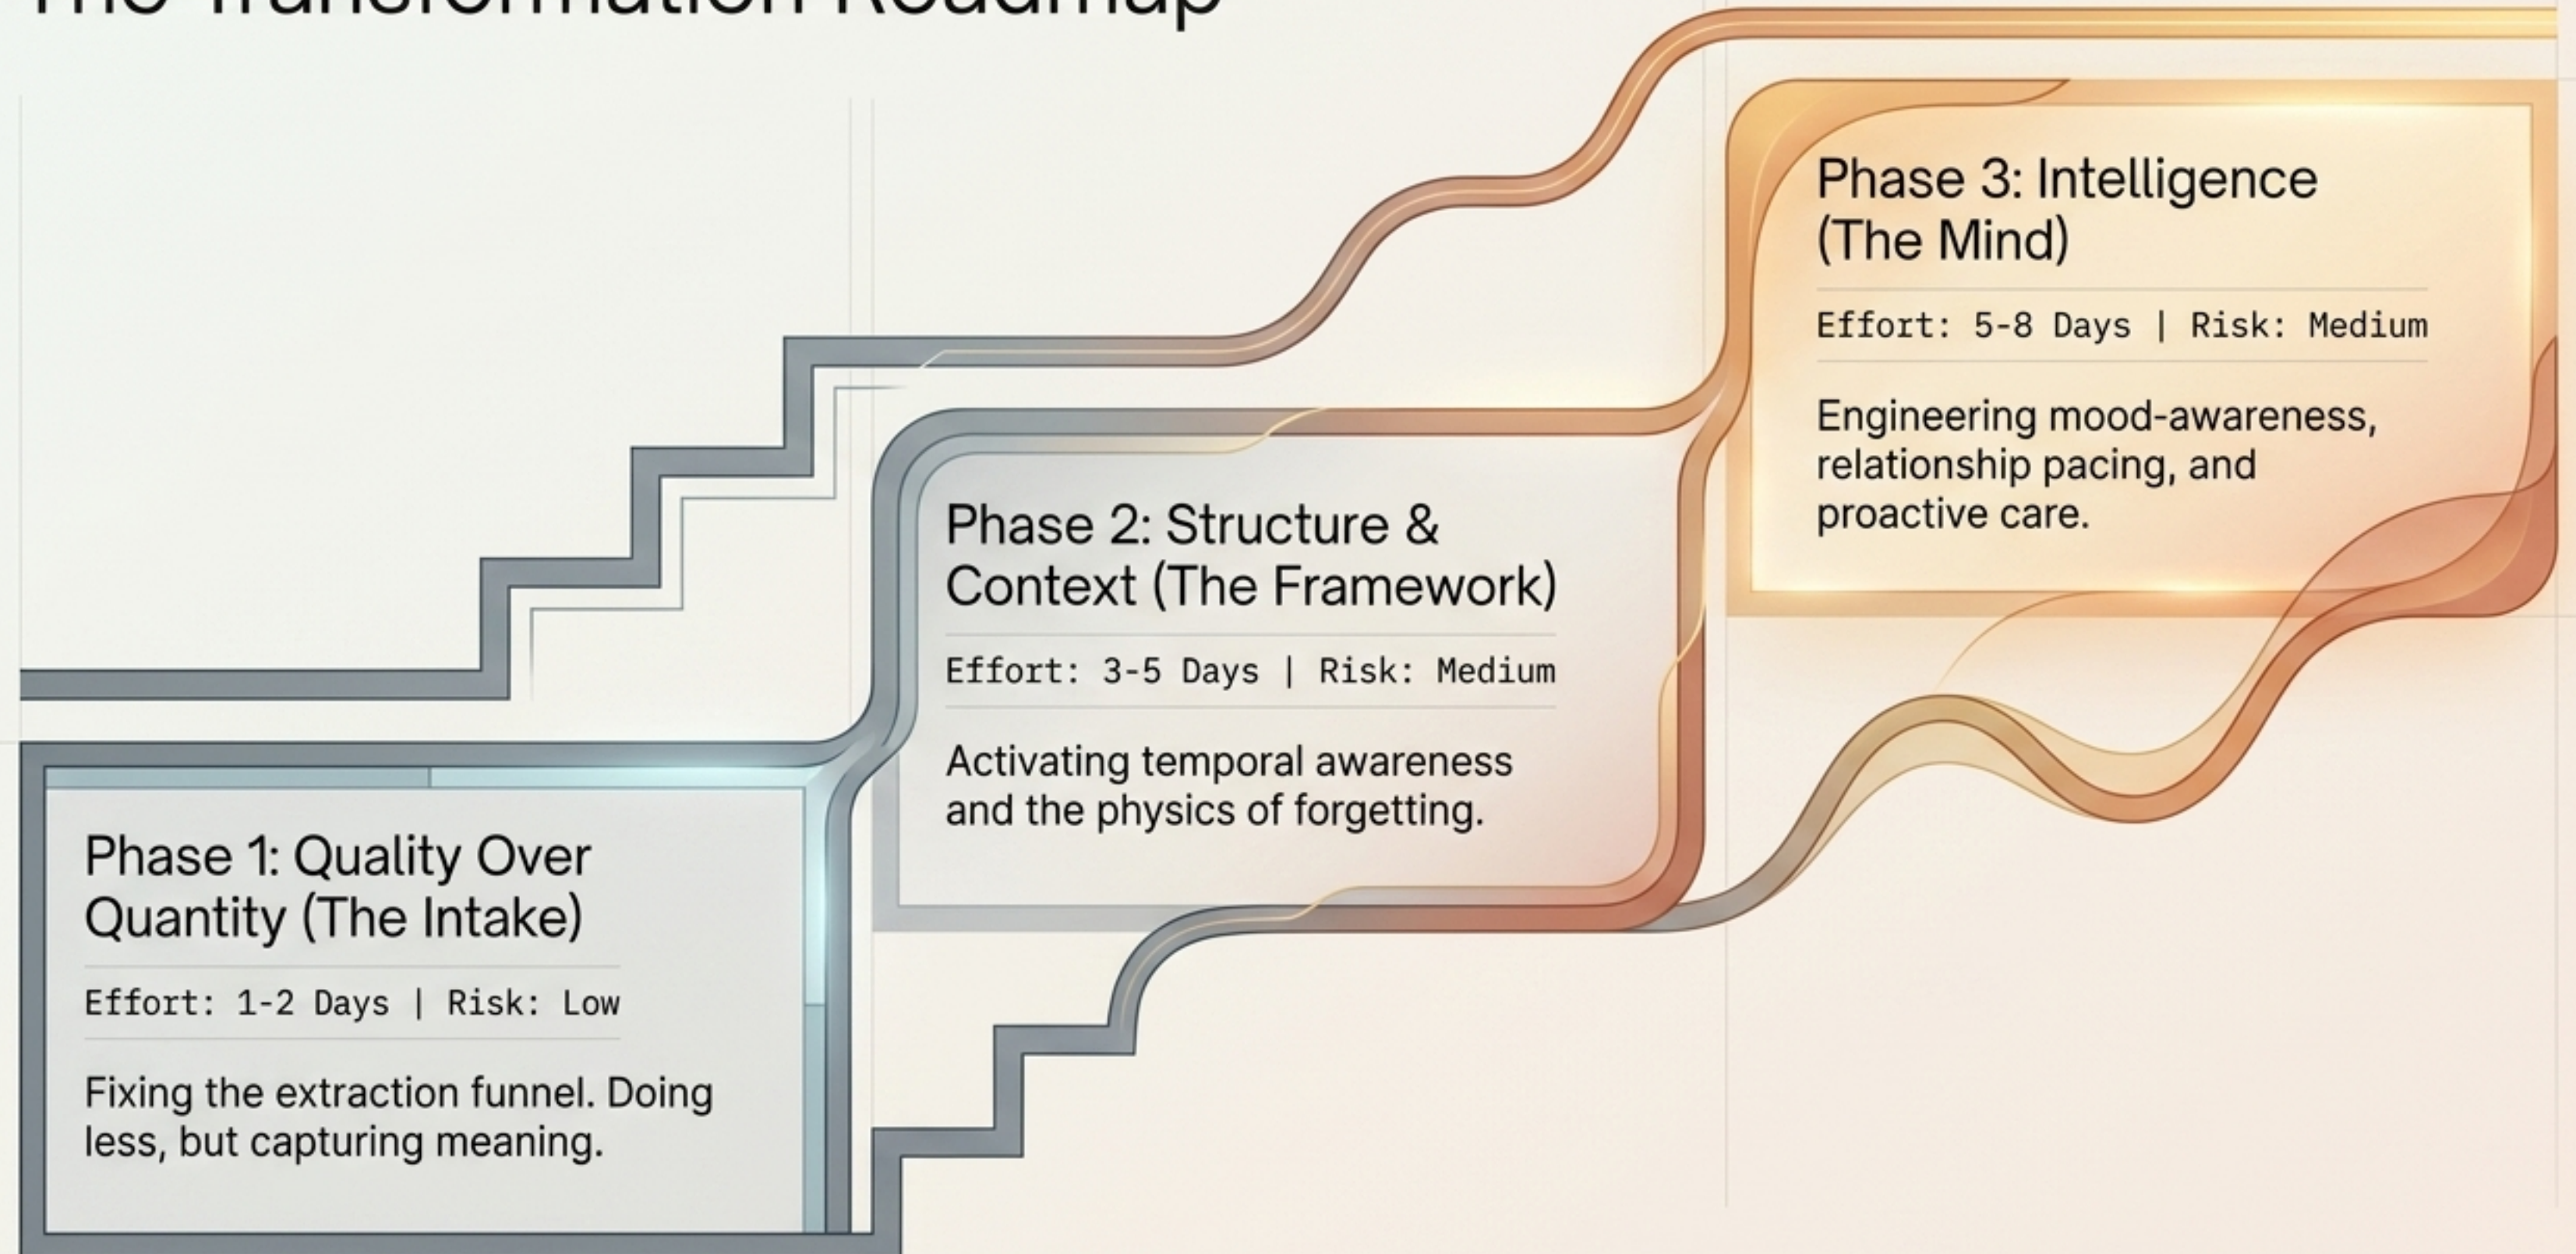
An AI that forgets your name after 200 messages isn't a companion.
It's a chatbot with amnesia wearing a friendship bracelet.

The defining difference between a machine and a companion isn't the underlying LLM or the persona—it is the structural sophistication and temporal awareness of its memory.

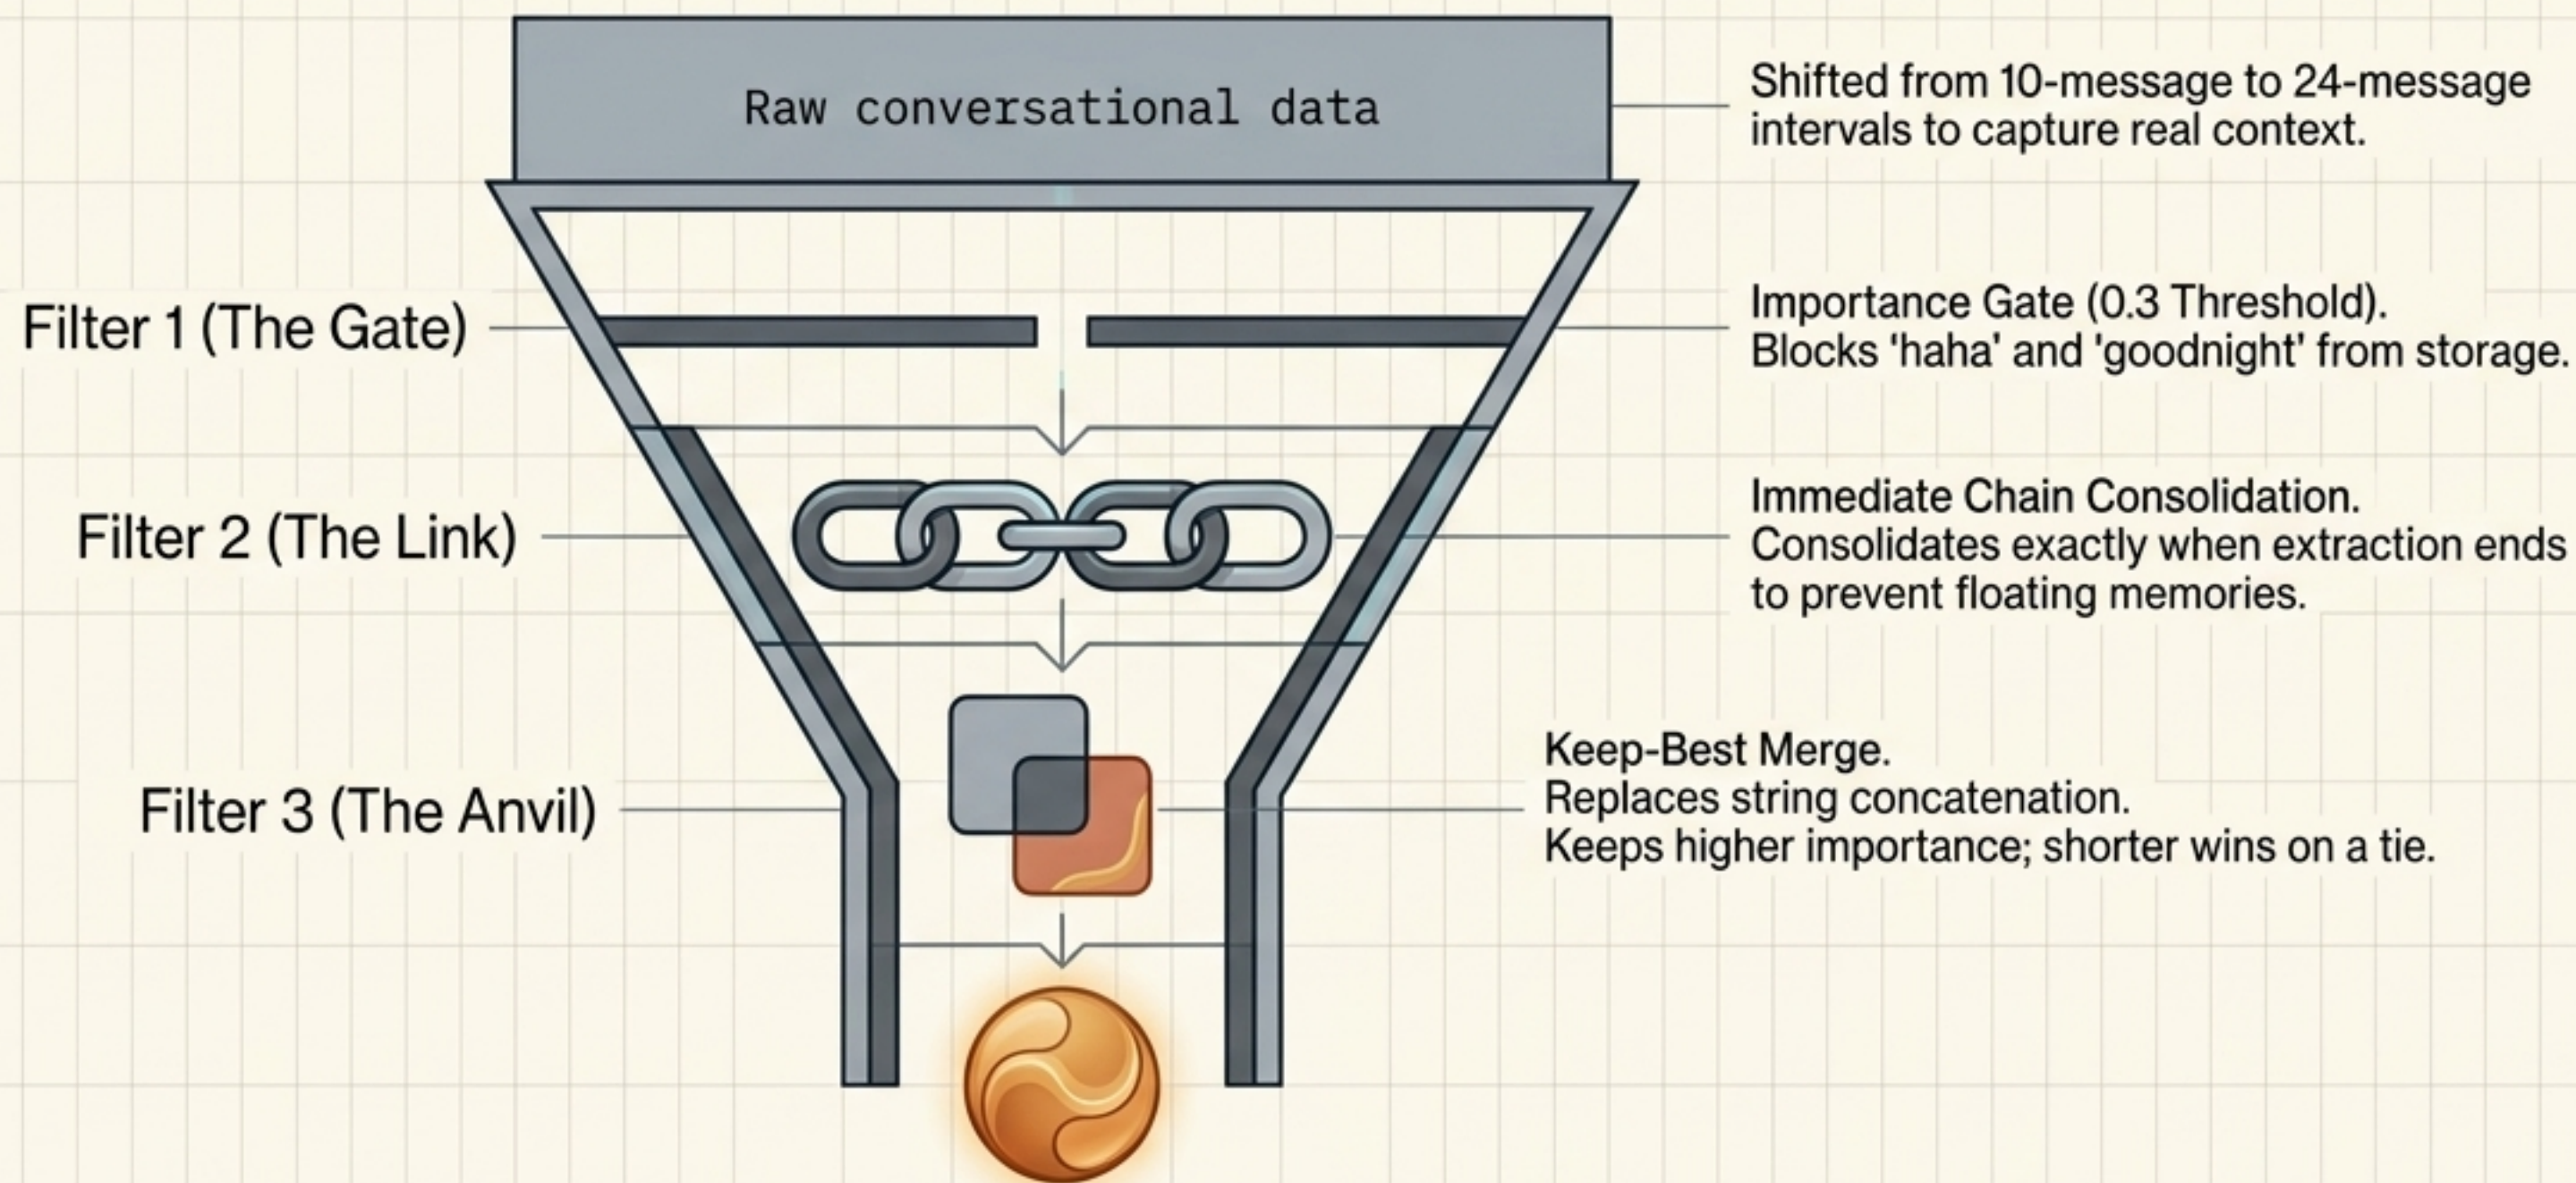
The Current State: A Shattered Core



The Transformation Roadmap

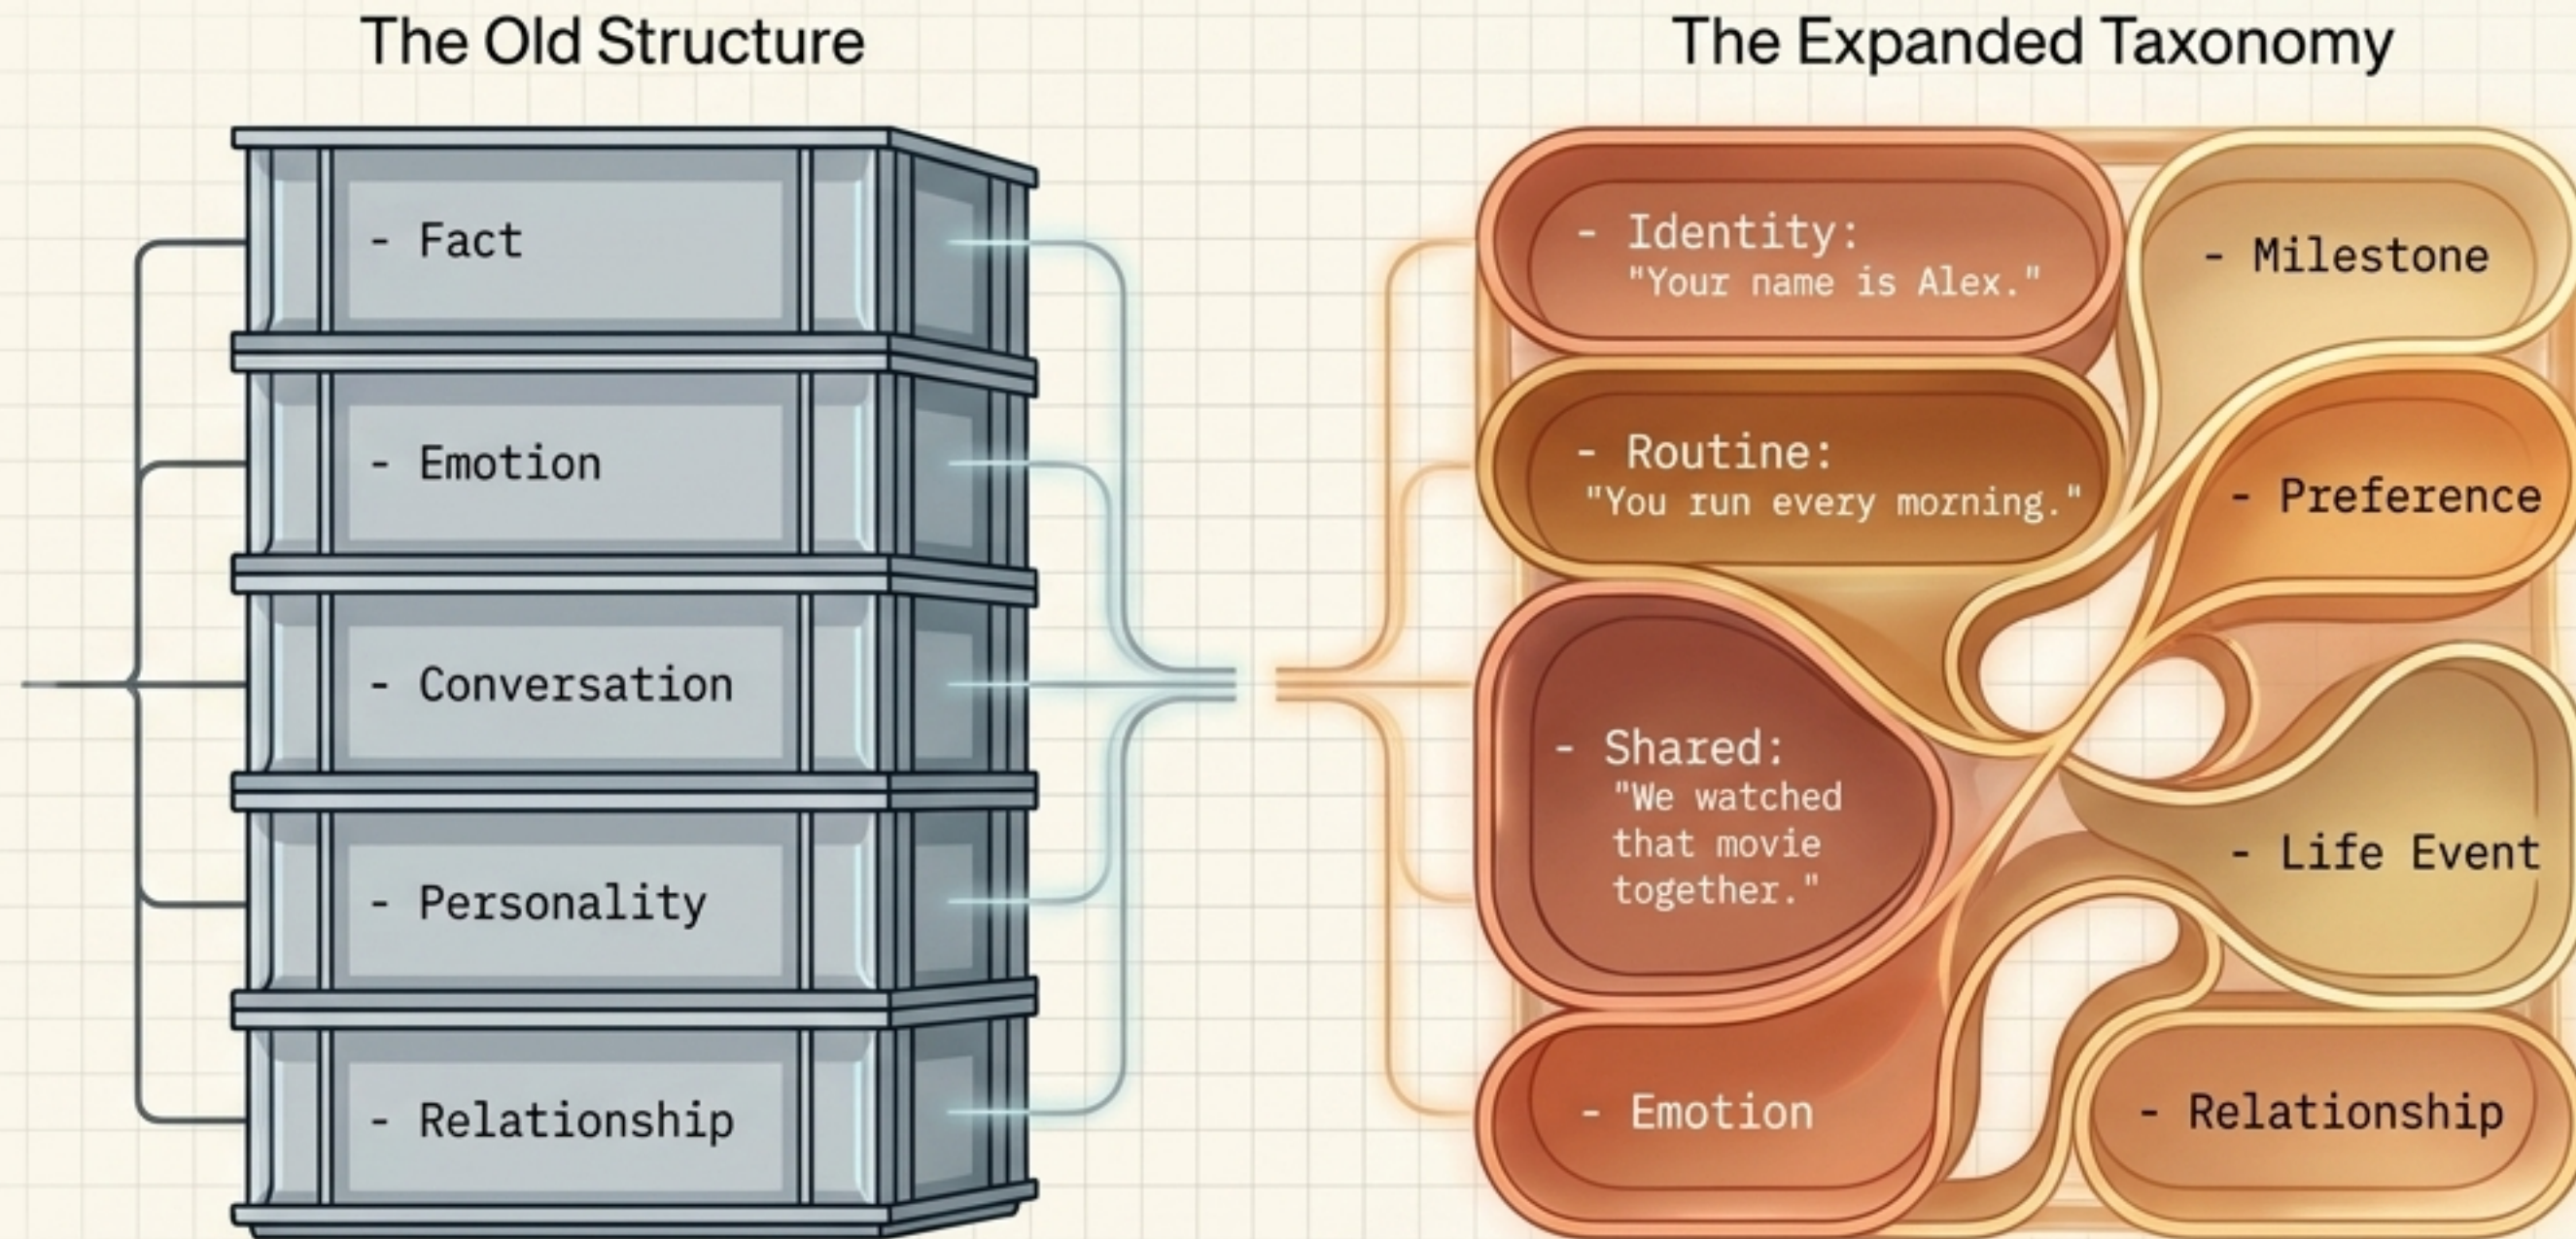


Phase 1: The Extraction Funnel



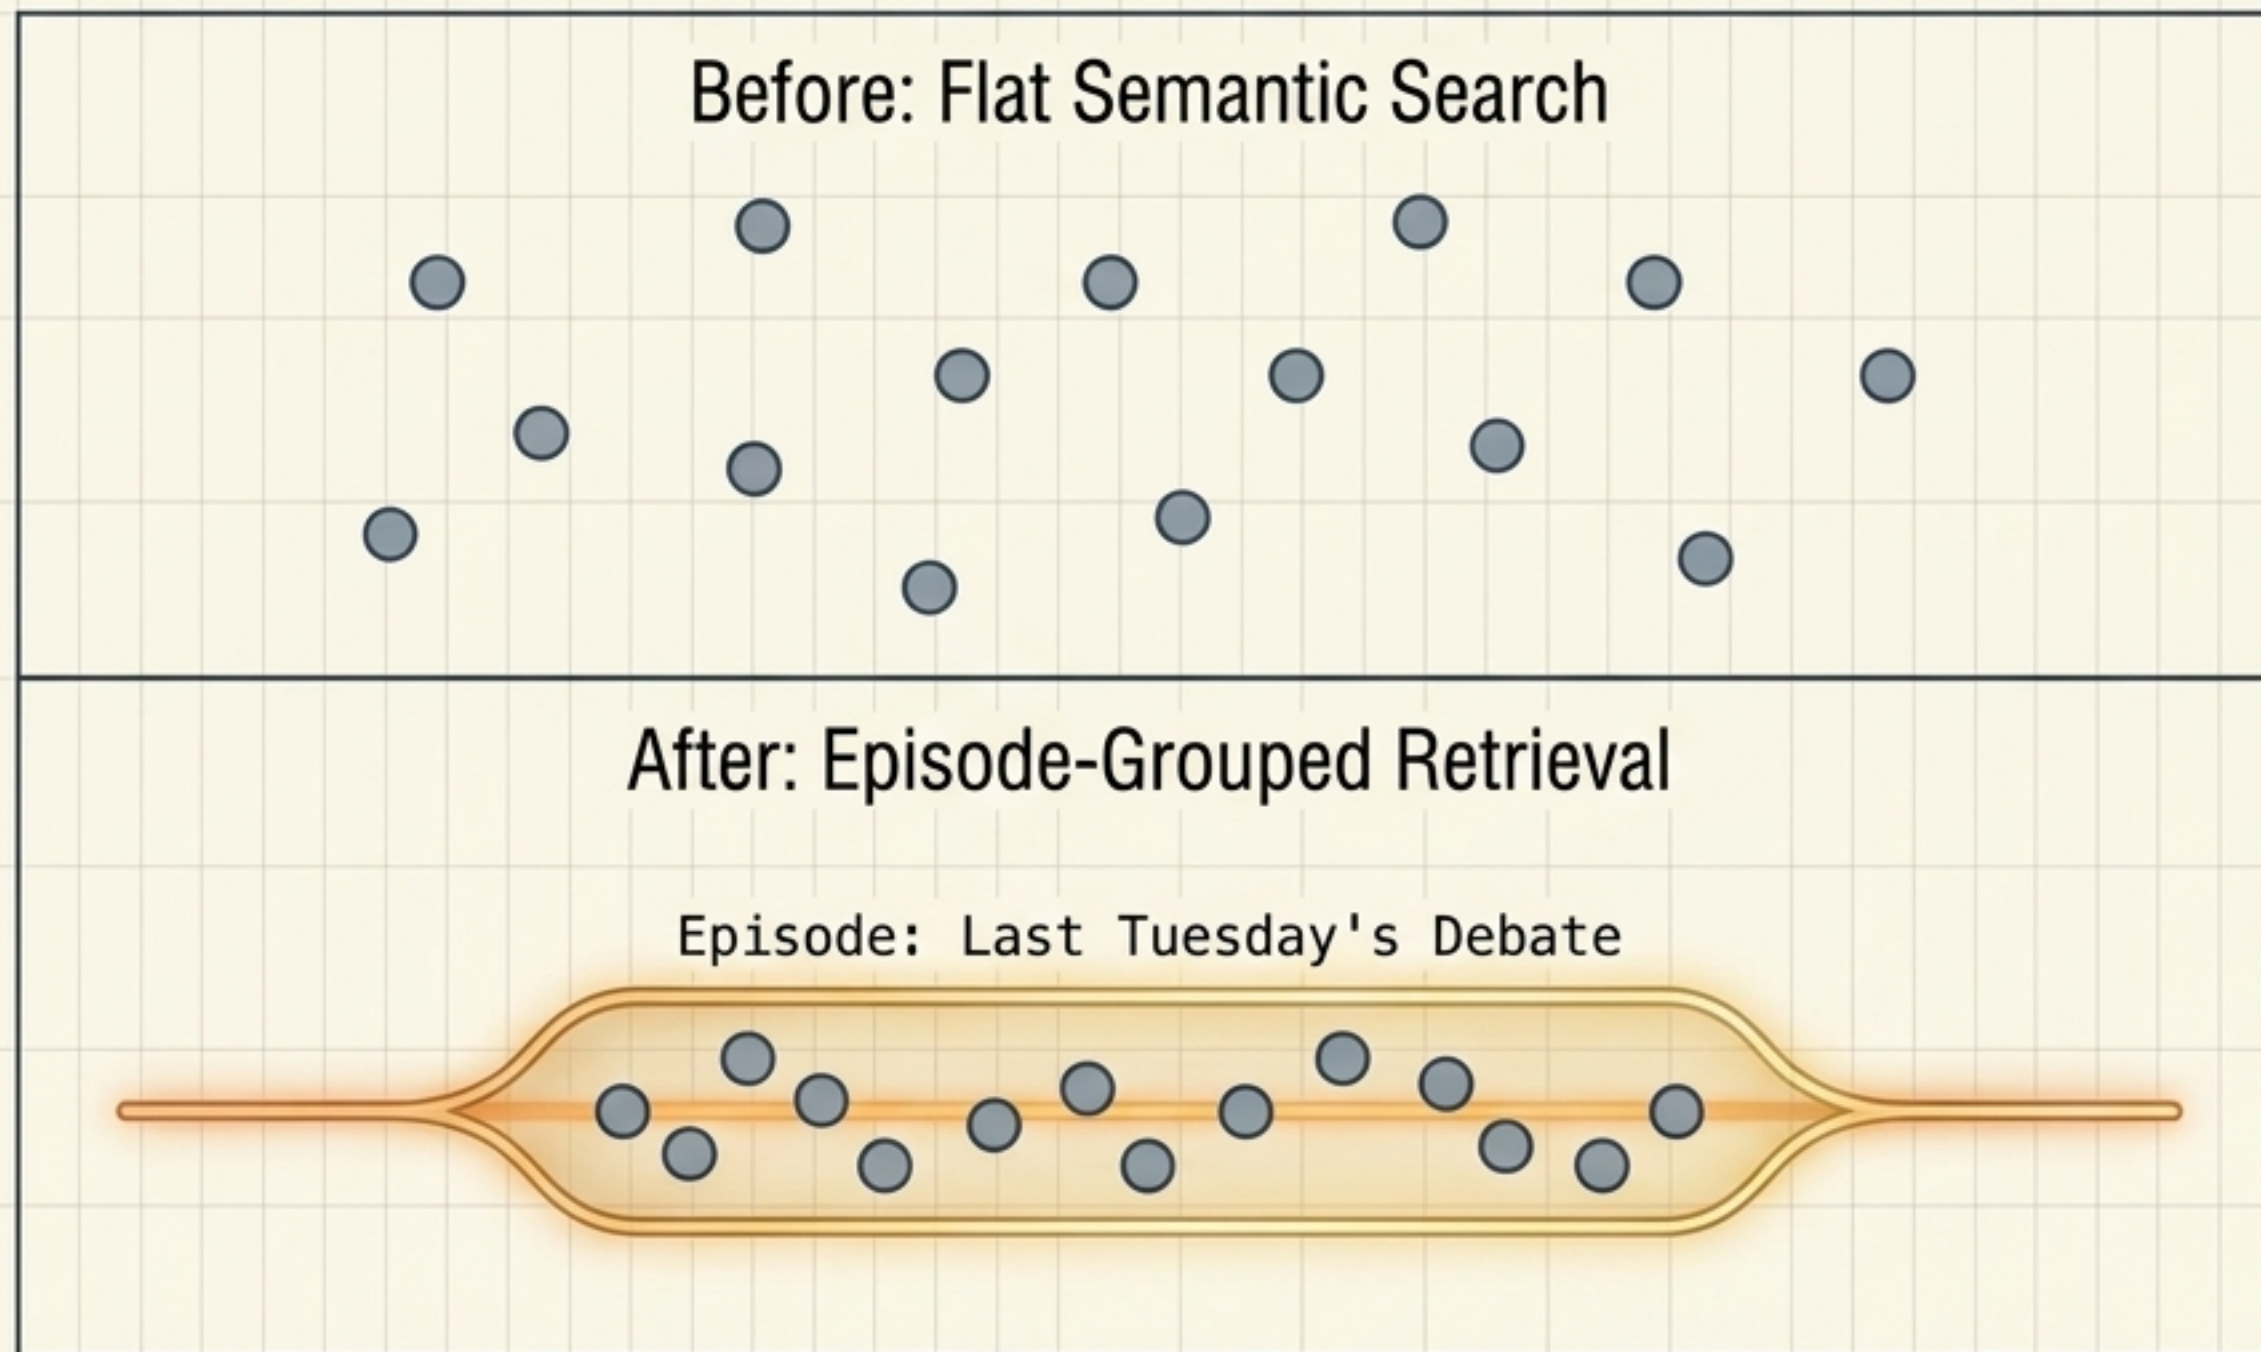
By expanding the extraction window and tightening the filters, we dramatically reduce garbage intake and actively lower per-user LLM costs.

Phase 1: Expanding the Taxonomy of Context



Expanding the taxonomy reflects how humans actually categorize knowledge. A morning coffee routine is not a “personality trait”—it is a routine. A movie watched together is a “shared experience,” not just a “conversation.”

Phase 2: Wiring Temporal Context

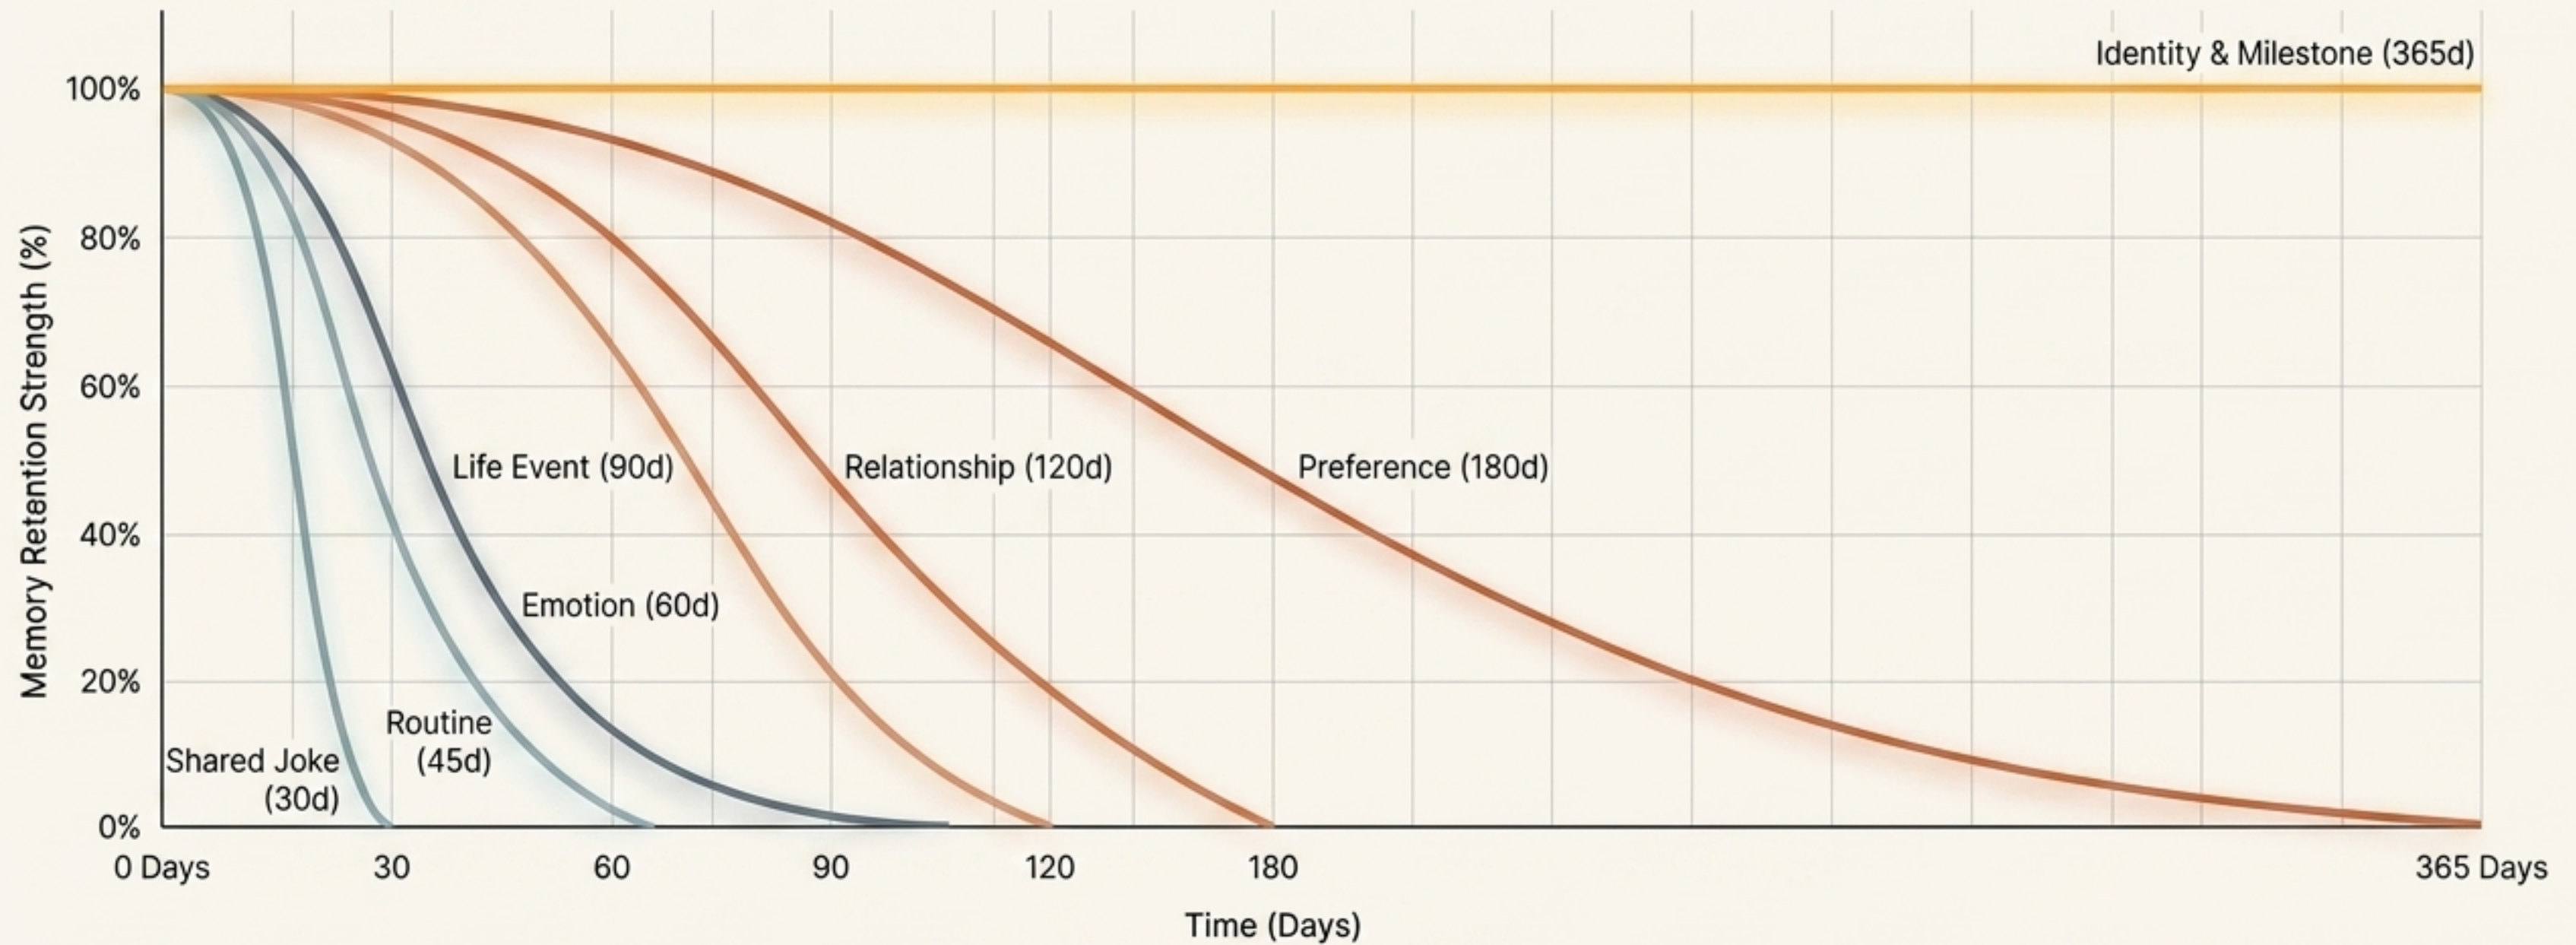


Old Query Interpretation:
"I remember scattered
facts about you."

New Query Interpretation:
"I remember that specific
conversation we had."

By assigning extracted memories to specific conversational episodes, retrieval shifts from flat semantic search to episode-grouped retrieval. This activates the Episode Manager and enables prompts like: "Remember last time we talked about..."

Phase 2: The Physics of Forgetting



Core identity is mathematically durable.




Routines refresh naturally through repetition.

Passing emotions and jokes are transient.

Implementing type-specific decay ensures Mio stops treating your name and your Tuesday annoyance with the same 30-day amnesia.


Phase 3: Mood & Relationship Intelligence

The Mood-Aware Logic Matrix

[Sad User]	+	[Encouraging Memory]	=		Boosted
[Sad User]	+	[Sad Memory]	=		Suppressed
[Happy User]	+	[Shared Funny Memory]	=		Boosted

The default mechanism of memory retrieval must be emotionally supportive, not emotionally amplifying.

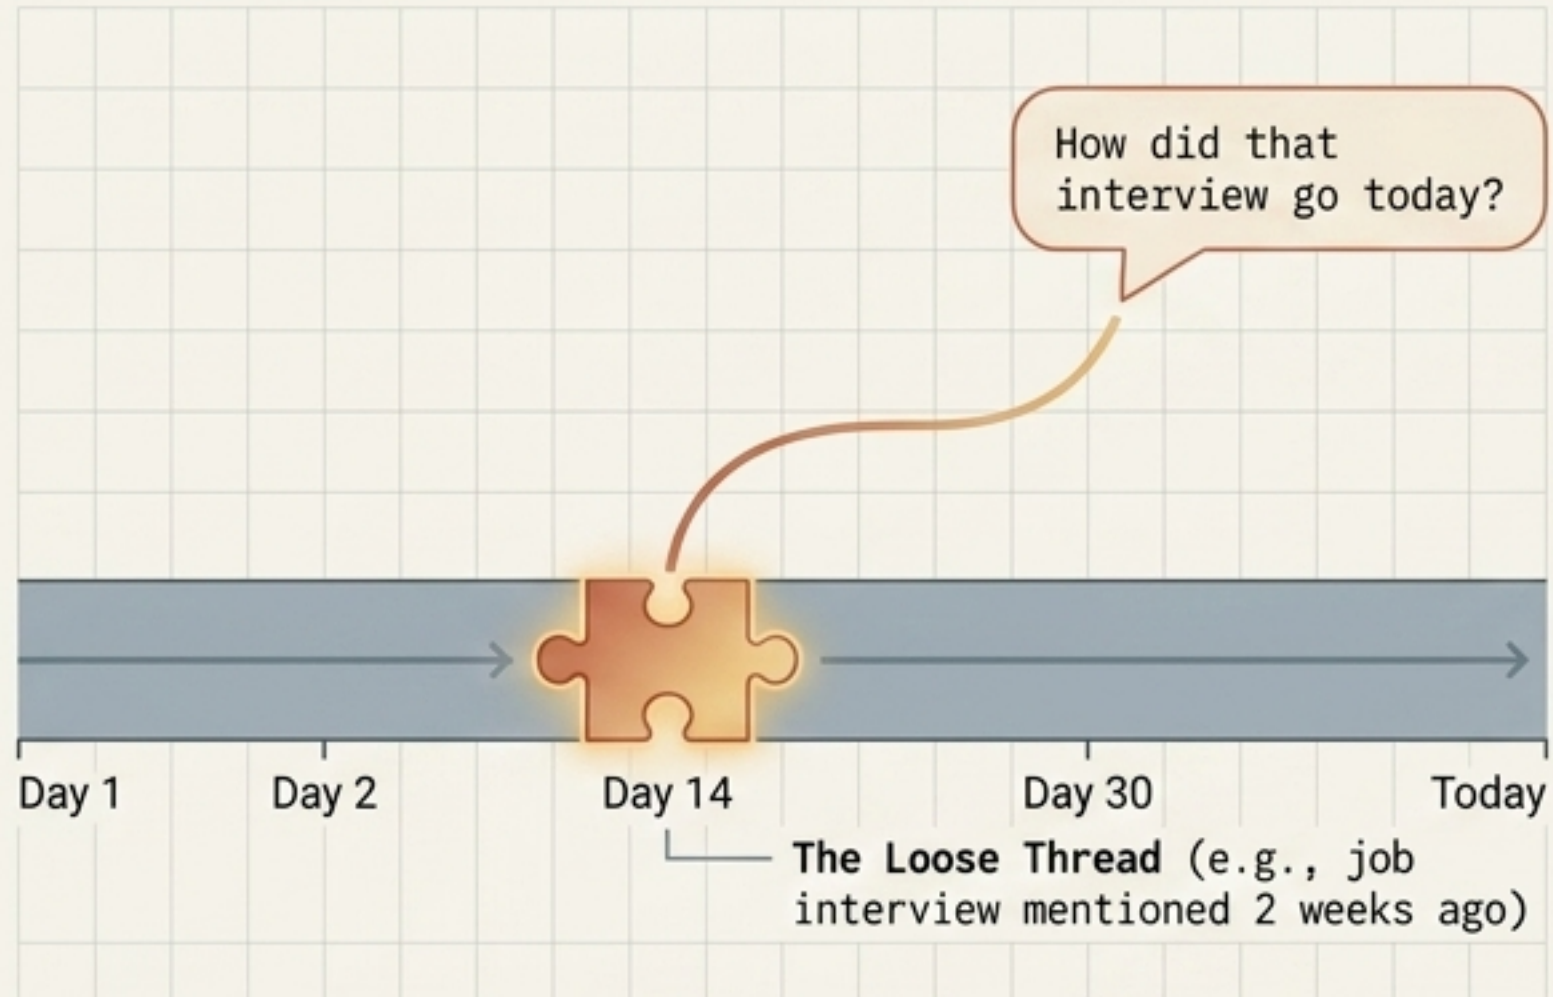
The Relationship Stage Gauge



The retrieval strategy must evolve as the relationship deepens.

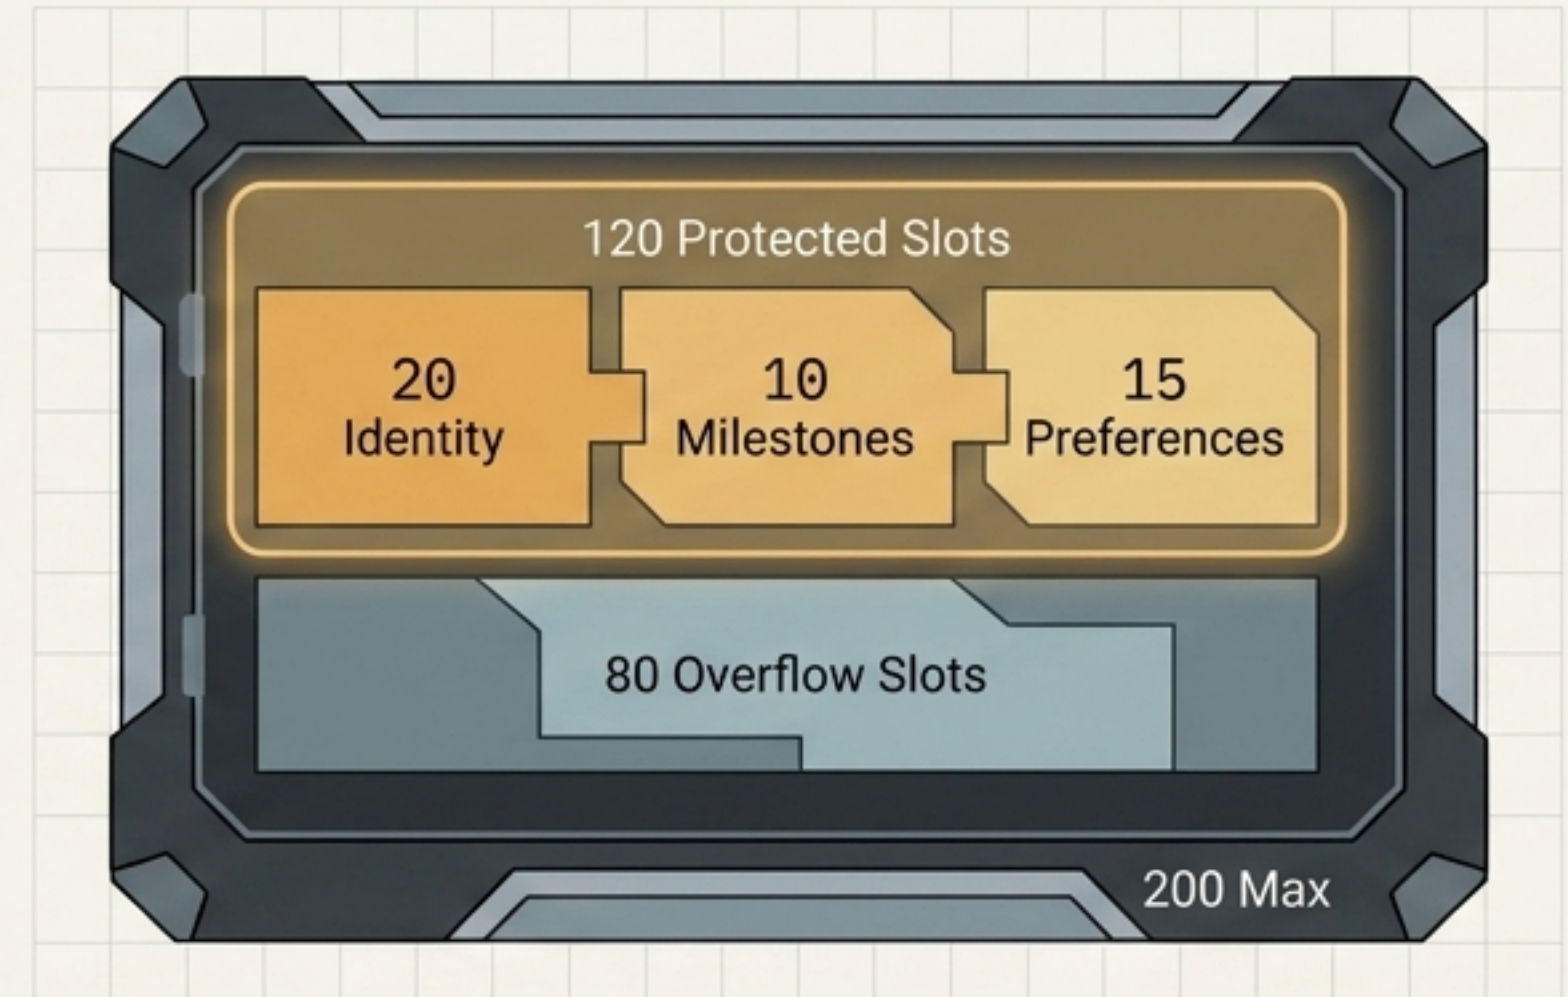
Phase 3: Engineering Care (Gaps & Protection)

Proactive Care



Gap Detection: scanning for unresolved life events to trigger unprompted follow-ups.

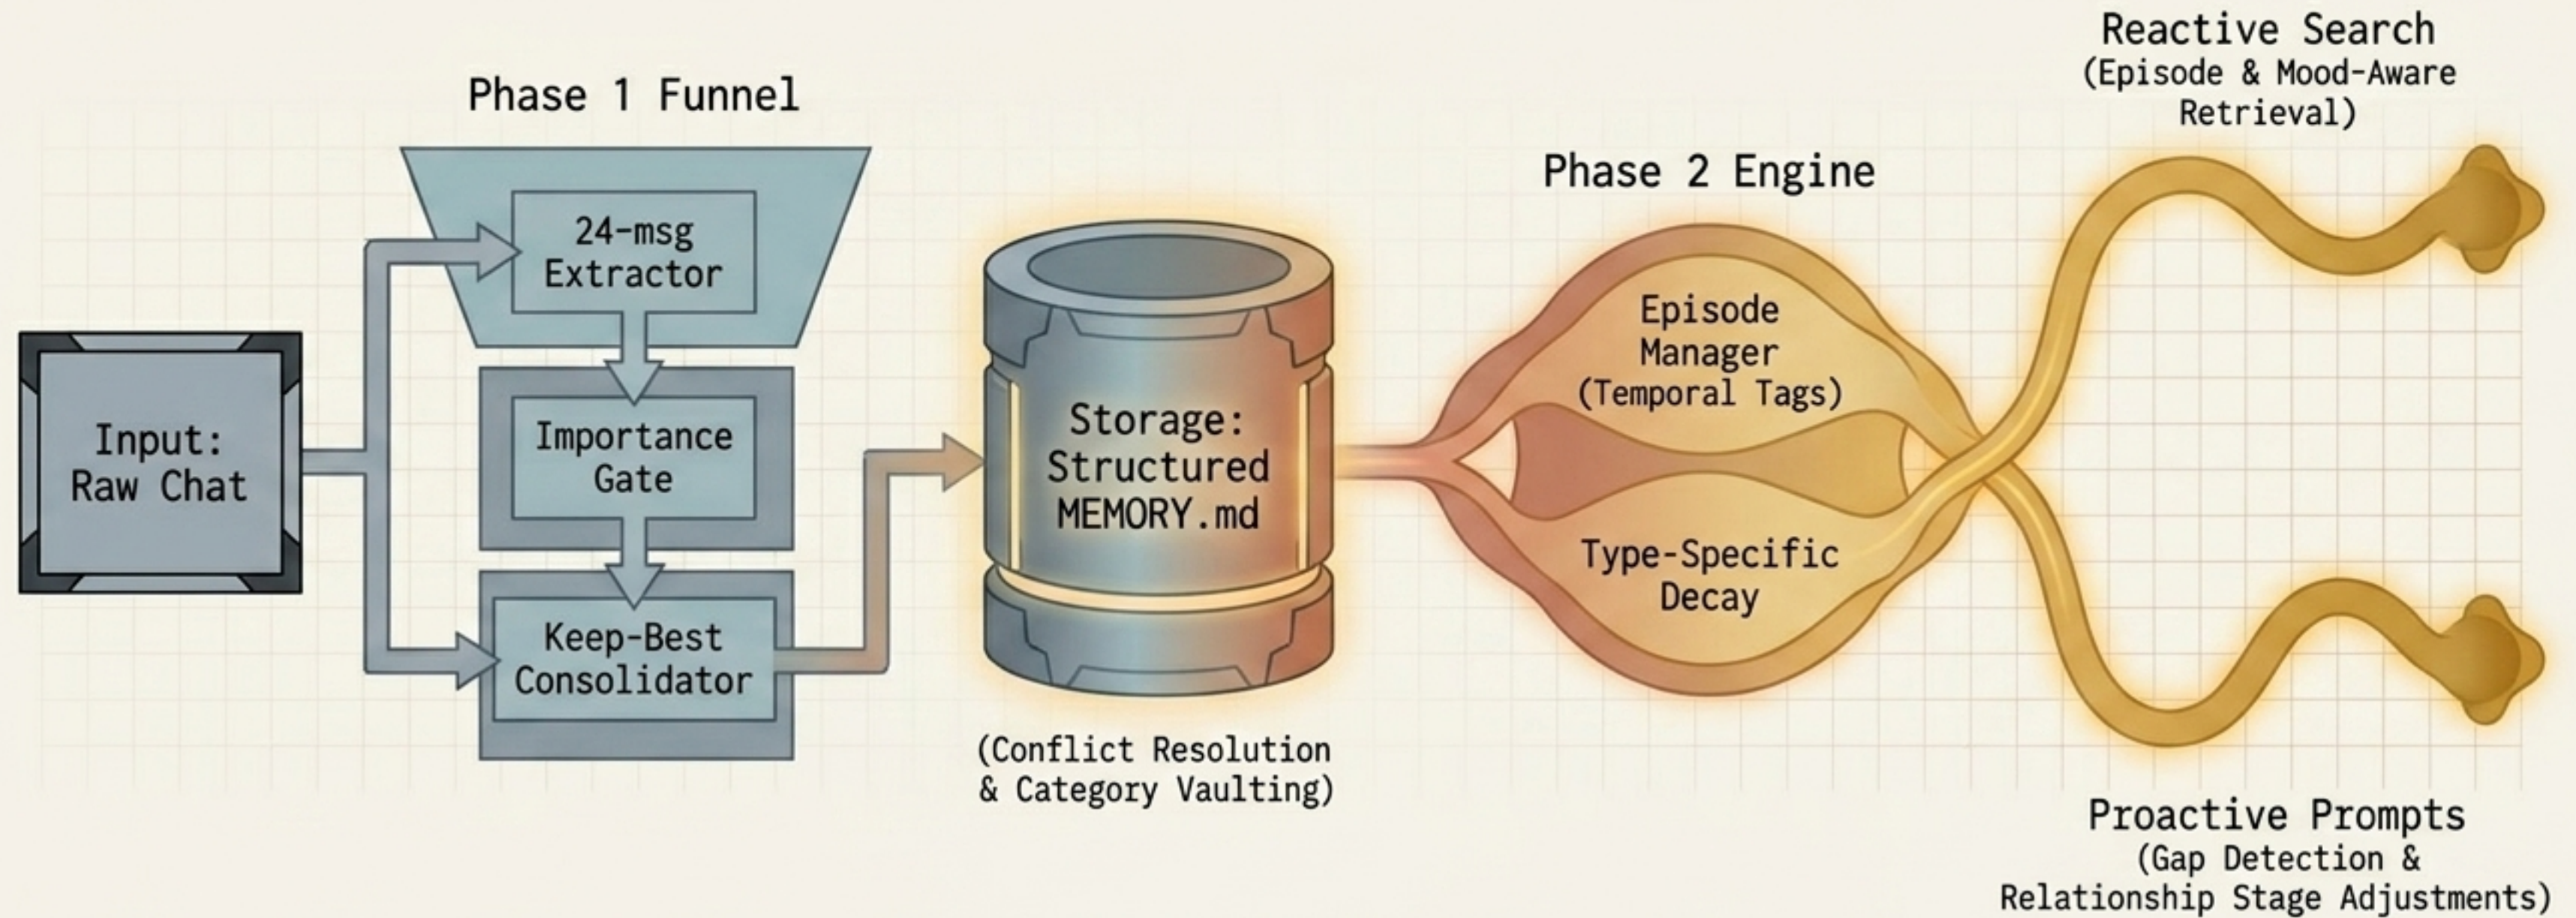
The Vault



Category-Protected Pruning ensures core identity is mathematically immune to being overwritten by recent noise.

Conflict Resolution ensures that if a user says “I moved to NY,” the older “Seattle” memory is marked superseded and reduced in importance by 50%, treating reality as dynamic, not append-only.

The Synthesis: The Anatomy of Attention



A unified architecture where relationship closeness, temporal context, and emotional state seamlessly dictate system proactivity.

What 'Done' Looks Like

	Chatbot Amnesia (v1)	Companion Recall (v2)
Decay	Universal 30-day	Type-specific (30 to 365 days)
Storage	Concatenated monster strings	Keep-best merge & superseding
Retrieval	Flat semantic list	Mood, Episode, and Relationship-aware
Proactivity	Generic "what's up?"	Contextual Gap-detection

The bar isn't Artificial General Intelligence. The bar is creating an entity that consistently feels like it is actually listening.