meko
By yugabyteDB

The Context Engine for the Agentic Enterprise

One Context. Many Agents. Governed Learning

Quality that scales with context: agents discover, humans govern, and everyone reuses, through a single MCP endpoint.

Meko is the universal, self-learning context engine for teams of agents and humans. It understands the structure inside your context and delivers it as efficient, governed infrastructure.

Every agentic application needs four things: a model to reason, context to draw on, tools to act with, and governance to keep it trustworthy. Models are easy to swap and tools are converging on MCP, but context is still built by hand, from scratch, for everyproject. Yet context decides whether everything else pays off: too little context and the model guesses; too much and therelevant facts get buried while cost climbs.

Meko breaks that paradox. It structures context automatically, tracks provenance from source to answer, and sends the model only the relevant slice for each turn, so quality stays high as context grows, at lower cost. It replaces the stitched stack of database, vector store, RAG pipeline, memory tool, conversation store, and artifact store with one serverless, multi-tenant system. Model-agnostic, delivered via MCP, REST, and UI.

Engineered Context

Small, High-Quality Context From a Large One

Meko structures your context, scores every search for relevance, and extracts the right subset for each turn. When recall is poor, it explains why and what is missing, then ingests data to close the gap automatically. Your agents get better answers with far fewer tokens,and smaller models can perform where larger ones were needed.

One Universal Context

Every Agent on the Same Page

The datapack holds all the context forone project, user, or tenant. Claude,Cursor, custom agents, and humans work from it at once, so what one agent learns is instantly available to another. Memory is private by default and shared by choice, and everything lives in one place instead of six silos.

Governed Learning

Agents Discover. Humans Govern.

An agent proposes a learning as private memory, an owner approves it with evidence attached, and only then does it become shared knowledge. One agent's opinion never silently becomes organizational truth. Decision traces show where every piece of context came from, who added it, and why.

93%
fewer input tokens: only the relevant subset of context goes to the LLM

~100%
accuracy in most runs, vs. ~75% with hand-rolled context

3–5x
cheaper tokens: engineered context lets smaller models perform

5+→1
systems consolidated: memory, vector DB, observability, glue

Figures are from Meko internal testing and pilot measurements; results vary by workload.

Everything You Need in a Single Context Engine

Context Engineering
Structures context and extracts only the relevant subset per turn, so quality holds as context grows.

One Universal Context
One datapack serves every agent, harness, and human, so no one starts from zero.

Governed Memory
Private-by-default memory with approval, promotion, editing, and deletion before it is shared.

Decision Traces
One audit surface for every memory read and write, promotion, and retrieval, with token tracking.

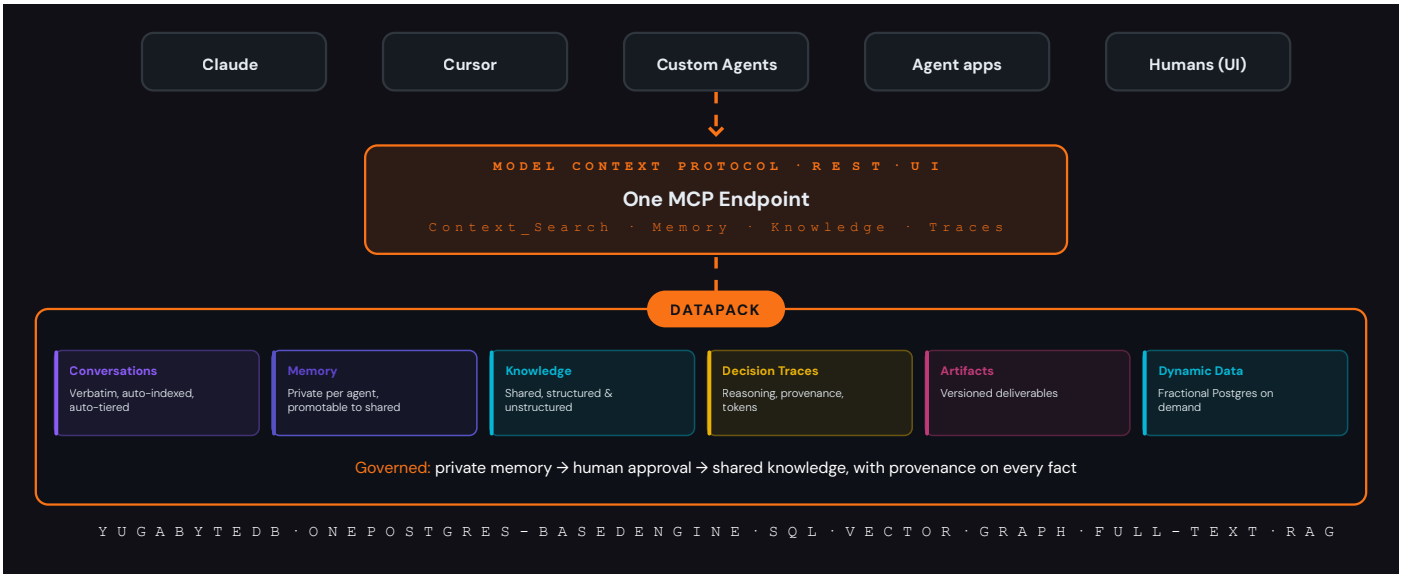
Self-Learning Loop
Every search is scored 0–1. Gaps becomeactionable needs and data is auto-ingested to lose them.

Single MCP Endpoint
Connect Claude, Cursor, and custom agents.context_search queries conversations, memory, and knowledge in parallel.

Serverless Architecture
Scale to millions of agents without complex infrastructure management requirements.

Serverless & Multi-Tenant
Millions of datapacks, one per project or user, with hot/cold auto-tiering and flat storage cost.

Model-Agnostic
Run on Meko Cloud or BYOC in your own cloud.Use any LLM, frontier or small.



Memory stores
MEM0 & CO.
Per-agent memory only. No sharing,governance, or audit.

Vector DBs / RAG
PINECONE, CUSTOM RAG
Context fragmented across 2–3 systems you stitch and sync yourself.

Observability
LANGSMITH, LANGFUSE
Trace what agents did, not what they learned or how it spread.

DIY on Postgres
POSTGRES + VECTOR +GLUE
Custom handoff schemas, extractioncode, and audit middleware.

Built on YugabyteDB
Meko's stateful layer runs on YugabyteDB, a horizontally scalable, PostgreSQL-compatible distributed database proven at mission-critical scale, with 5M+ clusters deployed across 100+ countries. Vector, SQL, graph, and search run in one engine, and each tenant's data is isolated in its own logical database. You get production resilience and the Postgres interface developers already know.

Who Meko is for
Teams putting agents into production: platform and data engineering leaders who own the agent stack, architects designing multi-agent systems, CTO and CIO organizations standing up a shared agent platform, and application owners in regulated industries where BYOC deployment and auditability are non-negotiable. The common thread is context that must outlive one agent and one session, governed across teams.

`Agentic software development` `Support for thousands of engineers` `Marketing campaigns: 3 months to hours` `Fraud detection`
`Solutions blueprints` `Compliance: EU AI Act, HIPAA, SOX, FedRAMP`

An agent discovers it. A human validates it. Meko remembers it and can explain it. Every future agent uses it. From isolated agent intelligence to governed institutional memory.

meko By yugabyteDB

MCP-native · Meko Cloud & BYOC

Get Started
mekodata.ai | mekodata.ai/discord