Your AI Agents, Supercharged

manus Dify aws LangChain LlamaIndex

What are the Data Challenges for AI Agents?

AI agent platforms don't scale like traditional apps; they can host millions of autonomous agents, each with its own state, memory, and evolving schema—effectively millions of database-like environments in parallel. Workloads swing from zero to extreme, mixing transactional bursts, analytical scans, and vector/semantic search as schemas change continuously. So the challenge isn't just "can it run," but whether it can run efficiently and affordably without metadata overload, tail-latency spikes, or runaway costs.

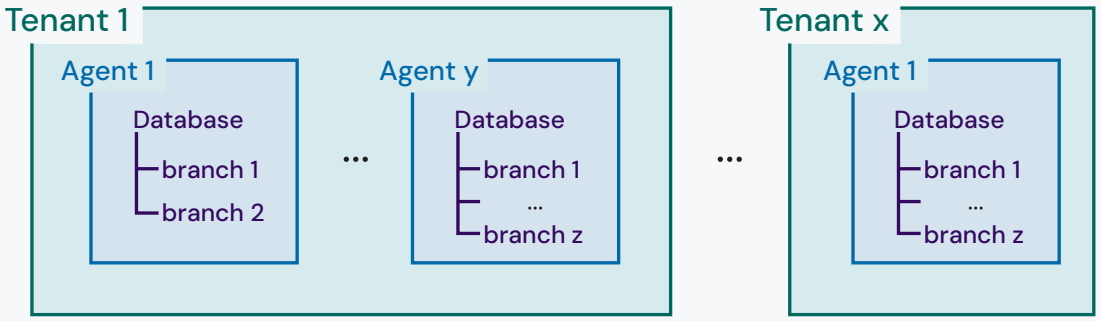

x × y × z tenants, agents, and branches multiply into millions of databases, exposing the scale and complexity of AI agent data challenges.

What are the Major Data Infrastructure Challenges in Building AI Agents?

- **Massive concurrency & volume:** Millions of agents reading/writing in parallel create petabytes of data and extreme concurrency that push past legacy limits.
- **Dynamic, per-agent schemas:** A platform may need to manage millions of unique tables, overwhelming metadata and isolation models.
- **Unpredictable workloads:** Mixed OLTP, analytics, and vector queries—changing as agents learn—demand robust ad-hoc performance.
- **Scaling vs. cost tradeoffs:** "DB-per-agent" explodes spend; "all-in-one DB" invites contention. Even serverless instances rack up idle cost.
- **Multi-modal needs:** Agents require vectors, JSON/knowledge, text search, and graph-like relations; spreading these across specialized stores and stitching them with ETL/ELT adds latency, failure points, and an ever-larger, harder-to-maintain stack.

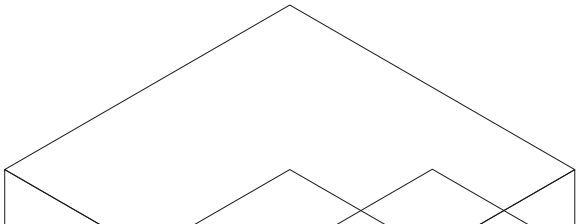

What is TiDB X? How Does it Solve Data Challenges for AI Agents?

TiDB X is a new, state-of-the-art TiDB architecture built for agent workloads that demand extreme scalability, elasticity, and intelligence. Its cloud-native architecture that separates compute from storage autoscales in real time and unifies transactional, analytical, and vector workloads. It delivers built-in vector search, hybrid retrieval, long-term memory, and LLM integration so agents can store context, reason over fresh data, and scale confidently.

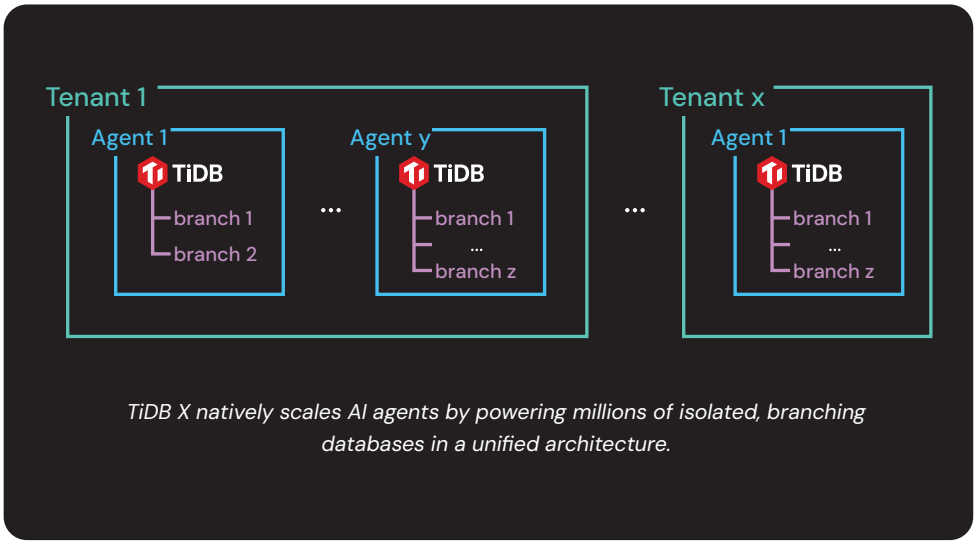

TiDB X natively scales AI agents by powering millions of isolated, branching databases in a unified architecture.

- **Extreme scale & concurrency:** Scale horizontally on demand to handle unpredictable spikes without degrading performance.
- **Multi-tenancy with isolation:** Efficient agent-level isolation via schema-level separation, resource groups, and copy-on-write branching enable safe parallel experiments and schema evolution without impacting neighbors.
- **Dynamic schema handling:** Support schema-per-agent patterns and rapid online evolution with minimal ops overhead.
- **Mixed-workload performance:** Optimize OLTP, analytics, and vector queries running simultaneously on the same platform.
- **Cost efficiency at scale:** Serverless, usage-based pricing scales resources up/down automatically, eliminating idle infrastructure cost.
- **Unified data types & fresh access:** Store rows, vectors, and semi-structured content together to keep knowledge current and consistent.

TiDB X + AI Agents: Key Use Cases

- Coding Agents
  - Personalized Website Builders
  - AI Dev Copilots
  - Multi-Repo Code Intelligence
- Agentic RAG
  - Dynamic Knowledge Assistants
  - Memory-Enhanced Chatbots
  - Multi-Source Reasoning Agents
- Workflow Automation Agents
  - AI Ops Orchestrators
  - Business Process Agents
  - Event-Driven Automation
- Enterprise Copilots and Analytics Agents
  - Internal Data Copilots
  - Customer Support Agents
  - Decision Intelligence Agents

Success Story: AI Agents Build Full-Stack Apps at Unprecedented Scale

manus

Manus 1.5 turns agents into autonomous full-stack builders that generate complete web apps—frontend, backend, auth, domain, analytics, and database—from a single prompt. Manus's agents have already created millions of database instances and branches to support independently evolving apps with highly unpredictable traffic and frequent schema changes. This scale is possible because TiDB X lets agents spin up isolated databases in seconds, branch them like Git for safe experimentation, and evolve schemas online without downtime.

TiDB X also powers Manus's hybrid workloads—transactional ops, real-time analytics, and AI-driven semantic queries—on the same data layer. With RU-based cost visibility and self-tuning signals, agents continuously optimize performance while keeping spend predictable—closing the loop between application evolution and infrastructure efficiency.

For more information, visit TiDB.io