InfoLibrarian™ ENTERPRISE AI · 2026 REALITY CHECK · KNOWLEDGE GRAPH THESIS

The AI Hype Cycle vs. the Cloud LLM Economy

Why slide-deck magic is stalling inside real enterprise data architectures — and why context, not more tokens, is the missing layer.

95% GenAI pilots with no measurable P&L

40%+ agentic projects cancelled by 2027

17% enterprises have deployed agents

93% orgs exceed their AI budgets

ADAPTED GARTNER HYPE CYCLE · GENAI + AGENTIC AI · MID-2026

Two technologies, two different points on the same curve

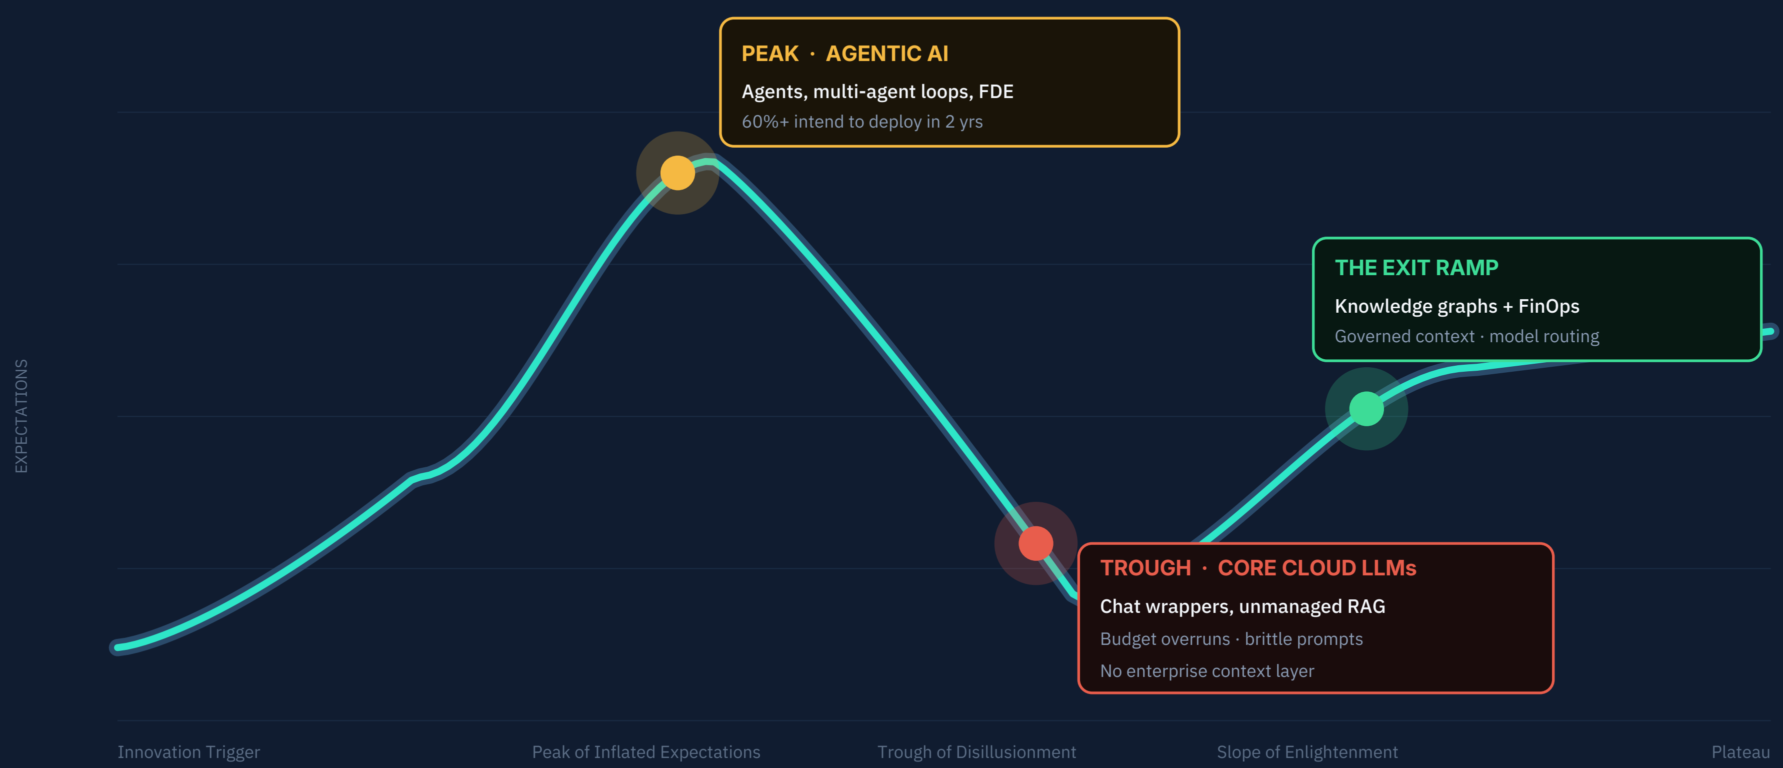

Core foundational models are sliding into the trough while agentic workflows are still cresting the peak. The collision is happening inside the enterprise stack.

THE HYPE PEAK

Agentic AI & custom wrappers

- **Autonomous agents**
  Promised end-to-end workflow automation with multi-agent loops that plan, tool-call, and act.
- **Forward-deployed squads**
  FDE teams wrapping flagship cloud models around each process — fast demos, fragile production.
- **Agent washing**
  Gartner: only ~130 of thousands of "agentic" vendors are the real thing. The rest is rebranded RPA.
- **Intention ≠ installation**
  17% deployed vs. 60%+ intending to. The installed base is years behind the narrative.

THE REALITY CHECK

Governance, cost, and data foundations

- **Token economics broke the business case**
  Agents resend long-lived context every step. Agentic tasks can burn ~1,000× the tokens of a chat turn. Price-per-token fell; total spend tripled.
- **No enterprise context layer**
  Cloud LLMs can read rows. They do not know how your company books revenue, who owns a policy, or which version is still valid.
- **Tech debt from bolted-on wrappers**
  Gartner: ≥50% of GenAI projects overrun budget from poor architecture. Most custom-model efforts get abandoned.
- **Governance is the hidden work**
  93% of enterprises hit permission/access issues. Teams spend ~77% of AI engineering hours on governance workarounds.

WHY THE CLOUD LLM ECONOMY IS NOT DELIVERING

The bill is a context bill — not a model bill

01 Stateless models, stateful business

Every turn re-ships history. Context rot and "lost in the middle" make long windows both expensive and inaccurate.

02 Vector RAG ≠ business understanding

Similarity search has no notion of ownership, authority, time, or permission. Demos work. Production does not.

03 Silos + over-sharing

Disconnected CRM, ERP, tickets, and docs. AI inherits years of access sprawl and invents confident answers across the gaps.

04 Wrong unit of value

Orgs bought tokens and copilots. Value only appears when workflows, cost-to-serve, and decisions actually change.

THE MISSING LAYER

Enterprise context via knowledge graph

A knowledge graph encodes entities, relationships, ownership, policy, and time — the map the LLM never had. Hybrid GraphRAG + governed retrieval is how slide-deck agents survive contact with a real architecture.

- **Grounded reasoning**
  Multi-hop answers over how the business actually works, not nearest-neighbor chunks.
- **Token discipline**
  Graph retrieval has cut token use 80–97% vs. naive RAG in published GraphRAG work.
- **Permissions & provenance**
  Who owns it, who can see it, which version is valid — encoded, not hoped for.
- **FinOps + routing**
  Smaller, cheaper models on a rich context layer beat flagship models on a pile of PDFs.

Sources: Gartner Hype Cycle for Agentic AI / GenAI 2026; Gartner cancellation forecast (40%+ agentic projects by EOY 2027); MIT Project NANDA (~95% of GenAI pilots, no P&L impact); McKinsey agentic economics (93% over AI budget); GraphRAG / enterprise KG literature. Figures are directional, not a substitute for primary research.

August 2026 · Directional figures from public analyst research · · www.infolibcorp.com