

Point of View

Financial Minefields of Agentic AI Adoption



Table of contents

1. Agentic AI cost risks	04
2. Key principles of agentic AI FinOps	06
3. Agentic AI FinOps Framework and best practices	07
▪ Pre-deployment cost intelligence	08
▪ Deployment: Orchestration dynamic cost control	08
▪ Post deployment – Sustaining performance and fiscal discipline	09
4. Conclusion: Agentic AI as strategic enabler	10



Agentic AI has moved from a theoretical concept to a boardroom priority. Poised to redefine automation, decision-making, and enterprise growth strategies, agentic AI systems are capable of autonomous reasoning, planning, and independent actions.

More than half (79%) of IT decision-makers agreed that agentic AI is one of their highest priorities over the next 12 months¹.

Unlike regular GenAI, which reacts to action prompts, agentic AI proactively interprets:



Context

Understanding user intent, time, and place, for example, knowing a task is urgent because it's end-of-quarter.



Behaviours

Recognizing user preferences and patterns, such as spotting when employees request reports or escalate issues.



Situations

Assessing the broader scenario or environment, like detecting system downtime or sudden changes in demand.

This helps formulate and execute next-step actions without human intervention. In fact, agentic AI actions can learn and adapt over time across multiple steps and services.

Enterprises are racing to adopt this game-changing variant of AI. Yet with unprecedented power comes complex and hidden business risks, especially regarding financial planning.

Unlike traditional cloud workloads, agentic AI systems are fast-changing and have highly distributed cost centres. New cost drivers now require an evolved approach to FinOps, including:

- ▶ **Model tokenization:** Breaking down data into smaller units (tokens) so AI models can efficiently process and understand information.
- ▶ **Continuous memory storage:** Persistently saving AI interactions and learnings so agents can recall and build on past experiences over time.
- ▶ **Agent orchestration loads:** The combined computing demand of coordinating multiple AI agents as they work together on complex tasks.
- ▶ **Human-in-loop dependencies:** The ongoing need for human oversight, validation, or intervention at various stages of AI system operations.

Today, agentic AI is at the centre of innovation. As the hype grows and adoption increases, enterprises sense the pitfalls of cost ahead in the journey.

Leading analysts predict that more than 40% of enterprise agentic AI projects will be shelved by 2027. This is typically due to escalating costs, unclear business value, and/or insufficient risk control. Hype initiatives often do not translate into genuine business value with accountable total cost of ownership (TCO).

Here’s why today, we explore why traditional FinOps frameworks fall short and introduce a practical, stage-wise approach to help you unlock business value while staying in control of costs and risks.

Key cost risks in agentic AI deployments

Despite enterprises recognising some visible costs, such as model licensing, computing, storage, data integration, and labour costs, the budget is often likely to go off track during project deployment due to the poorly understood, less visible costs.

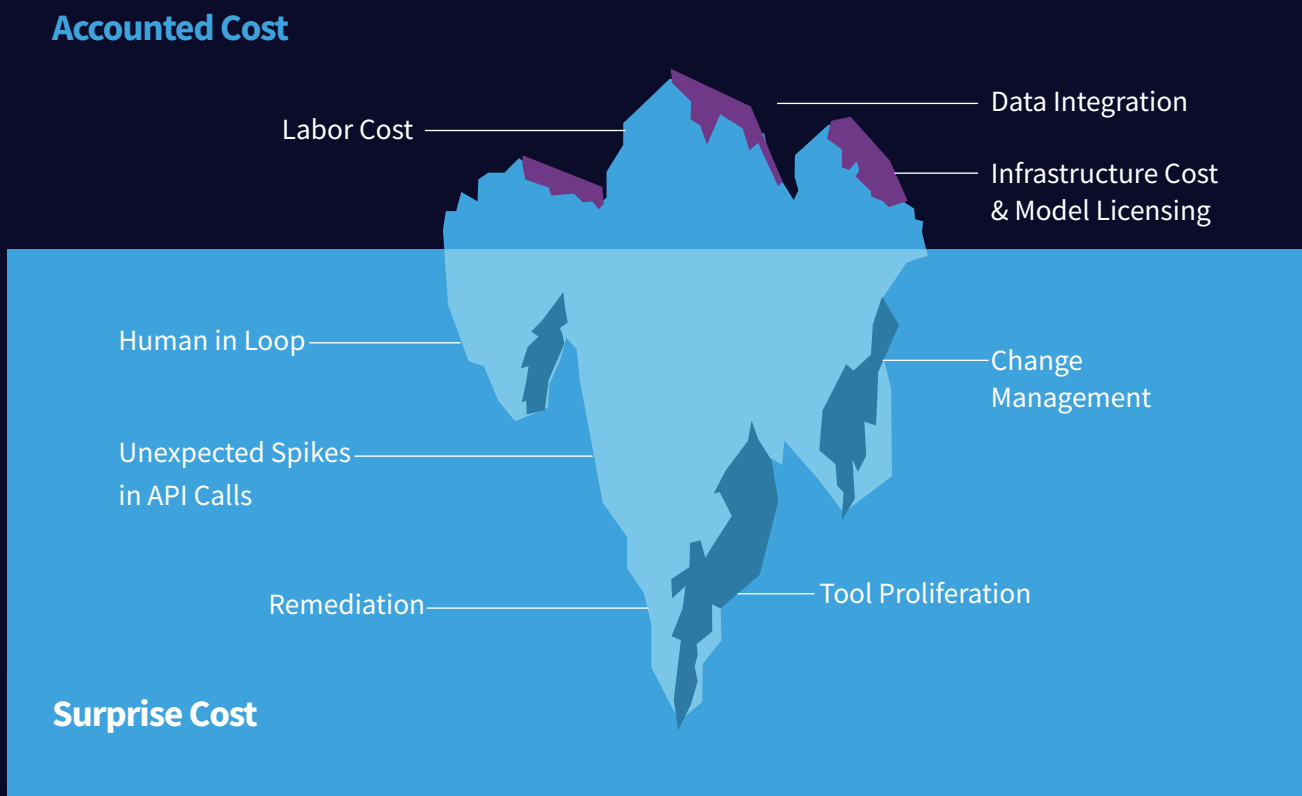


Figure 1: Cost risks of agentic AI



Human-in-the-loop

Even the most promising AI agents can not eliminate the need for significant human involvement during the project life cycle and post-deployment.

Process architects, data stewards, AI ethics, security, and ops teams must oversee, validate, and maintain agentic AI systems. A certain percentage of AI output will require human validation in many products. This is a fundamental truth.



Change management

Building an agentic product is one thing. However, integrating it into the system is a mammoth task. Integrating agentic systems demands reimagining workflow, team education, and cultural resistance. All of these require significant transformation costs, which can be more than the cost of the product itself.



Unpredictable spikes in API calls

Agentic AI requires continuous communication between multiple complex systems. A prototype may not accurately reflect their business-as-usual (BAU) intensity.

Scaling an agentic project in production can lead to API call spikes, inflated data transfer fees, storage costs, and network usage. Multi-agent collaboration can exponentially compound resource usage and, therefore, costs.



Remediation

There is no 100% perfect system. Agentic AIs can become faulty and overreactive, wasting associated costs related to models, API calls, and transactions.

A comprehensive understanding of potential cost incidents is needed to develop and implement appropriate guardrails.



Tool proliferation

Enterprises actively engaged in R&D often subscribe to multiple AI tools. Usually, this represents capability redundancy and overcomplexity.

The fully loaded costs of autonomy and scalability frequently overlook these expenses, which can negatively affect perceived ROI.



Security and compliance

Agentic AI is likely not currently a mature capability in any enterprise, and enterprises are not fully prepared for it yet.

While full autonomy is a north star vision, new risk vectors in data privacy, access control, and the right security model require investment in categorising data, monitoring, and auditing to avoid expensive breaches and regulatory penalties.

Core principles for agentic AI FinOps

The dynamic, decentralised, and autonomous nature of agentic AI requires an evolution of the traditional FinOps frameworks and processes. Reactive approaches of manual cost review and post-factum optimisation are simply not sustainable for scalable agentic AI.

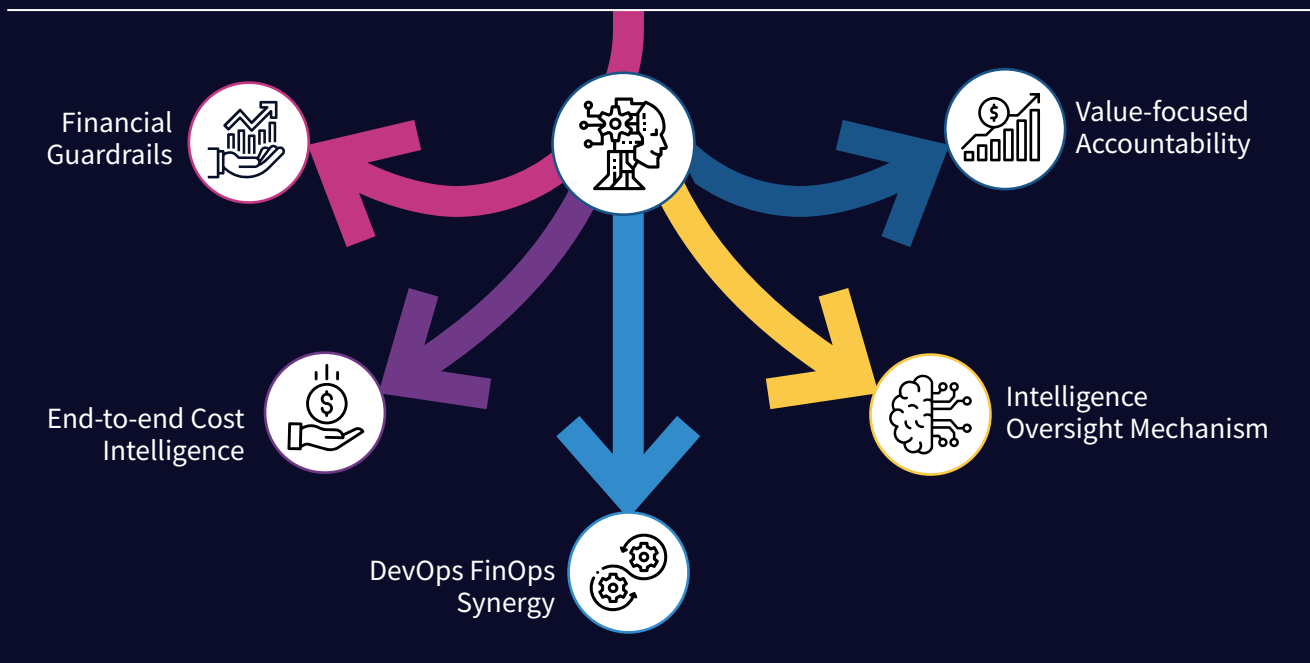


Figure 2: Key principles of agentic AI FinOps



End-to-end cost intelligence

AI-powered FinOps tools monitor agentic AI workloads to look for deployment inefficiencies, usage anomalies, and other opportunities for savings around the clock. These tools can identify and predict cost-optimisation opportunities, formulate solutions, and intelligently deploy them.

The system can automate optimizations like right-sizing resources, discarding unused assets, or halting cost-anomalous processes for human review.



Finance guardrails

Deploying agentic AI becomes far less risky when architects embed cost controls by design. Budget-aware logic, rate-limiting, and model tiering can avoid unpredictable expenditure.



DevOps-FinOps synergy

Tight integration of CI/CD and operations pipelines can embed agentic FinOps principles by enabling real-time financial visibility and agility across development and cloud engineering workflows.



Real-time reporting or value-focused accountability

Real-time dashboards can track agentic spend and value KPIs, ensuring that business cases and success measures are transparent and actionable.



Intelligent oversight mechanisms

While the system can maximise financial automations, it must rely on reactive human oversight, review, approval workflows, and other policy controls to proactive autonomous enforcement. This prevents financial risks and meets regulatory requirements.

Agentic AI FinOps framework

To mitigate financial risks and drive business value, teams must weave FinOps principles into each stage of development.

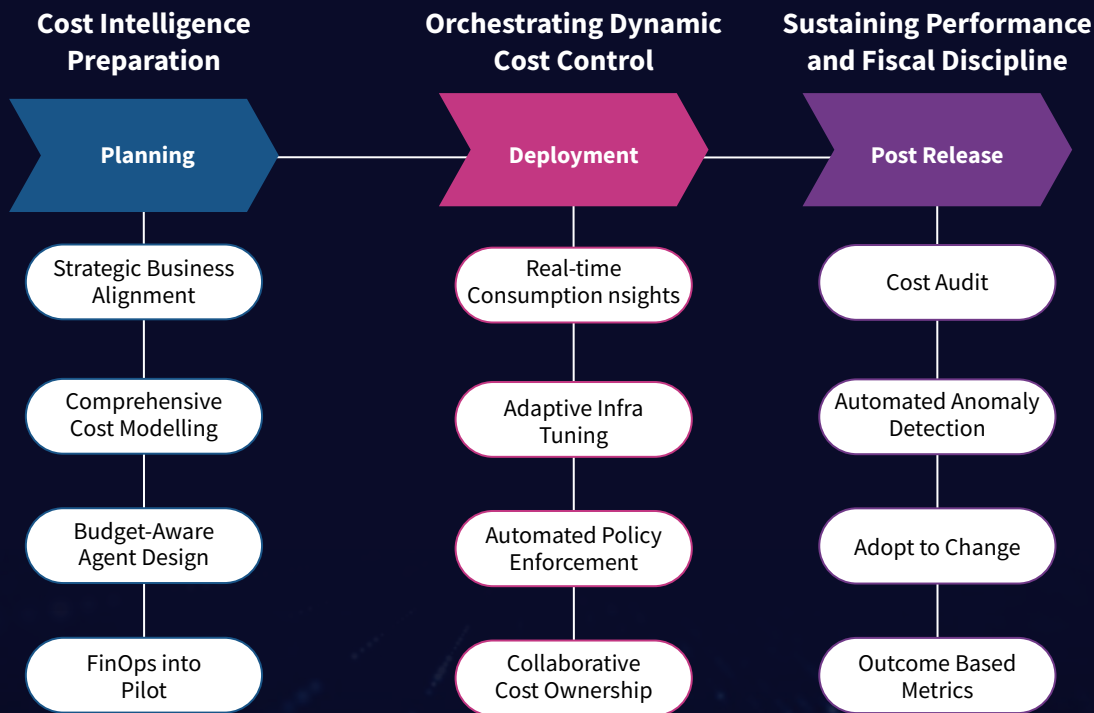


Figure 3: Agentic FinOps framework across the development lifecycle

Pre-deployment cost intelligence

- Align each agentic AI deployment with a distinct business objective, not merely a technical one.
- Set clear measures of application success and plan your approach based on a realistic understanding of agentic AI capability.
- Build a cost model that looks beyond foundational Cloud infrastructure like compute, storage, and licensing. Ensure it covers cost and hours:
 - To improve the quality and integration of your data
 - Required human involvement
 - Manage change and handle integration complexity
 - To cover unexpected usage spikes and the cost of course corrections
 - To implement the security controls
- Design agents with budgets in mind, using features like rate-limiting, smart model selection, and dynamic switching to strike the right balance between performance and cost.
- Enforce cost review from the early stages of pilot programs and ensure practices are followed before final approval.

Deployment

Orchestration dynamic cost control

Set up real-time agentic FinOps dashboards using observability tools to track costs where it matters most, including:

- API calls
- Memory usage
- LLM token consumption
- Orchestration overhead

Tune your infrastructure on the fly by:

- Routing workloads intelligently based on real-time demand
- Using dynamic rightsizing and serverless (pay-as-you-use) models to support flexible, cost-effective scaling

Put automated cost control policies in place by:

- Enforcing mandatory policy standards
- Applying consistent tagging practices
- Detecting and remediating cost anomalies as soon as they appear

Make cost management a team initiative by:

- Bringing finance, engineering, and business teams into continuous alignment
- Ensuring everyone has a clear line of sight into spend versus value at every stage

Post-deployment

Sustaining performance and fiscal discipline



Keep the improvement cycle alive by:

- Investigating what's really driving spend
- Optimizing agent behaviours for better efficiency
- Refining how you allocate costs
- Documenting lessons and sharing them across future projects

Build agentic systems that are flexible enough to evolve alongside changes in AI tech, compliance rules, and pricing models.

Measure what matters. Use smart dashboards to track technology-led business outcomes like request processing time, decision accuracy, and revenue growth.

Conclusion

Agentic AI as strategic enabler

Agentic AI represents a significant leap on the innovation scale. Enterprises are racing towards this opportunity and have already started prototyping and pilot programs.

Travel, education, hospitality, banking, retail, and other domains are recognising the possibilities that agentic AI brings to their operations. We are witnessing a technology driving a “new normal” for digital enterprises, likely beginning to challenge the market competitiveness of currently lagging organizations.

However, the promised benefits can erode if teams fail to conduct proper financial stewardship due diligence. Agentic FinOps about cost-cutting, maximising the return, and accounting for the value of every dollar invested.

As AI Agents are introduced into critical business processes, building cost-aware processes in every stage of the agentic AI development lifecycle, backend by FinOps automation and human governance will be decisive in determining which organisations are really leading and sustaining innovation.

Call To Action

Agentic AI is undoubtedly reshaping business operations and interaction with systems. With complex systems and ever-growing foundational models, managing financial risks requires a methodical approach. By adopting a comprehensive agentic FinOps framework at every stage of these intelligent agent deployments' lifecycle, enterprises can remain agile, resilient, and focused on being market leaders without being hammered by cloud costs.

References

1. *Techstrong.ai, Survey Surfaces Surge in Agentic AI Investments in the Enterprise, Mike Vizard, March 10, 2025, <https://techstrong.ai/agentic-ai/survey-surfaces-surge-in-agentic-ai-investments-in-the-enterprise/>*
2. *Gartner, Over 40% of Agentic AI Projects Will Be Canceled by End of 2027, June 25, 2025, <https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-a-gentic-ai-projects-will-be-canceled-by-end-of-2027>*

About the authors



Ashutosh Dixit

Principal Director of Cloud & Infra Consulting, LTIMindtree

Ashutosh, a passionate technology evangelist and AWS Ambassador for LTIMINDTREE, serves as a Principal Director at LTIMindtree. He leads the Cloud Strategy, Advisory, and Consulting team. He provides CIO advisory services worldwide and is an experienced solutions architect with expertise across all hyper-scalers, FinOps, security, and adoption frameworks. A renowned thought leader, he often speaks at industry events.



Pandiya Kumar Rajamony

EVP, Cognitive Infrastructure Services

Pandiya Kumar Rajamony (Pandi) serves as the Executive Vice President of Cloud and Infrastructure Services at LTIMindtree, bringing over 30 years of global expertise in technology leadership. Known for driving cloud and AI-powered transformations, Pandi specializes in scaling enterprise technology, fostering innovation, and crafting sustainable digital strategies. Passionate about enabling businesses to thrive, he focuses on building future-ready technology ecosystems leveraging modern cloud architectures, advanced data platforms, and AI solutions.

LTIMindtree is a global technology consulting and digital solutions company that enables enterprises across industries to reimagine business models, accelerate innovation, and maximize growth by harnessing digital technologies. As a digital transformation partner to more than 700 clients, LTIMindtree brings extensive domain and technology expertise to help drive superior competitive differentiation, customer experiences, and business outcomes in a converging world. Powered by 83,000+ talented and entrepreneurial professionals across more than 40 countries, LTIMindtree — a Larsen & Toubro Group company — solves the most complex business challenges and delivers transformation at scale. For more information, please visit <https://www.ltimindtree.com/>.