



# Going for the Gold with FinOps Forward and AI

**Jevin Jensen**

Research VP, IDC

April 2026





# Agenda

**1. Opportunities, Spending, and Enterprise Budgets for 2026**

**2. Industry and Market Overview**

**3. Maturing FinOps**

**4. AI for FinOps**

**5. FinOps for AI**

**6. AI Agent Future**

**7. Conclusion**



# Opportunities to grab the Gold

Industry and Market Overview



## Hybrid cloud still the king

Adding capabilities to manage **both** public and private cloud is necessary as the majority of enterprises still leverage this architecture.



## AI costs

A lack of forecasting and understanding of AI costs impacts the TCO of projects. Accurate cost calculator of future AI costs and best practices needed.

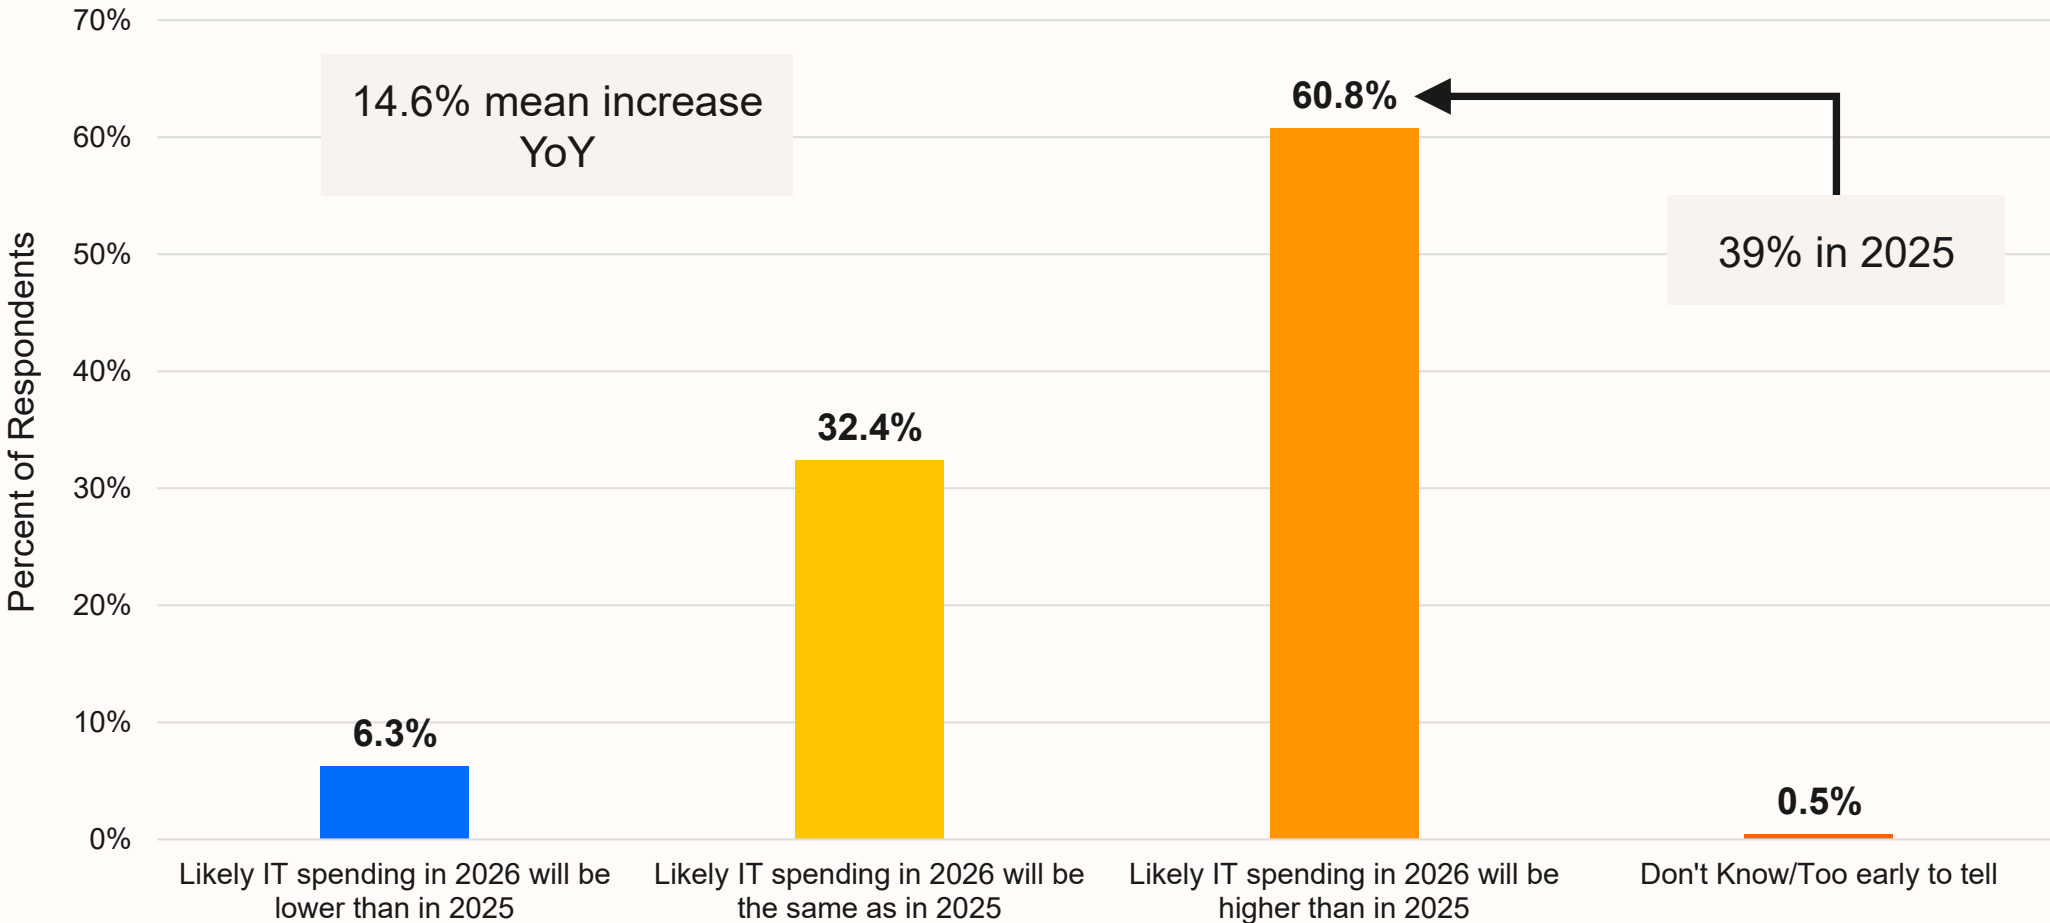


## FinOps maturity

Enterprises struggle to move to FinOps "run" maturity and maintain it. Professional/managed services and easier-to-use products may help.

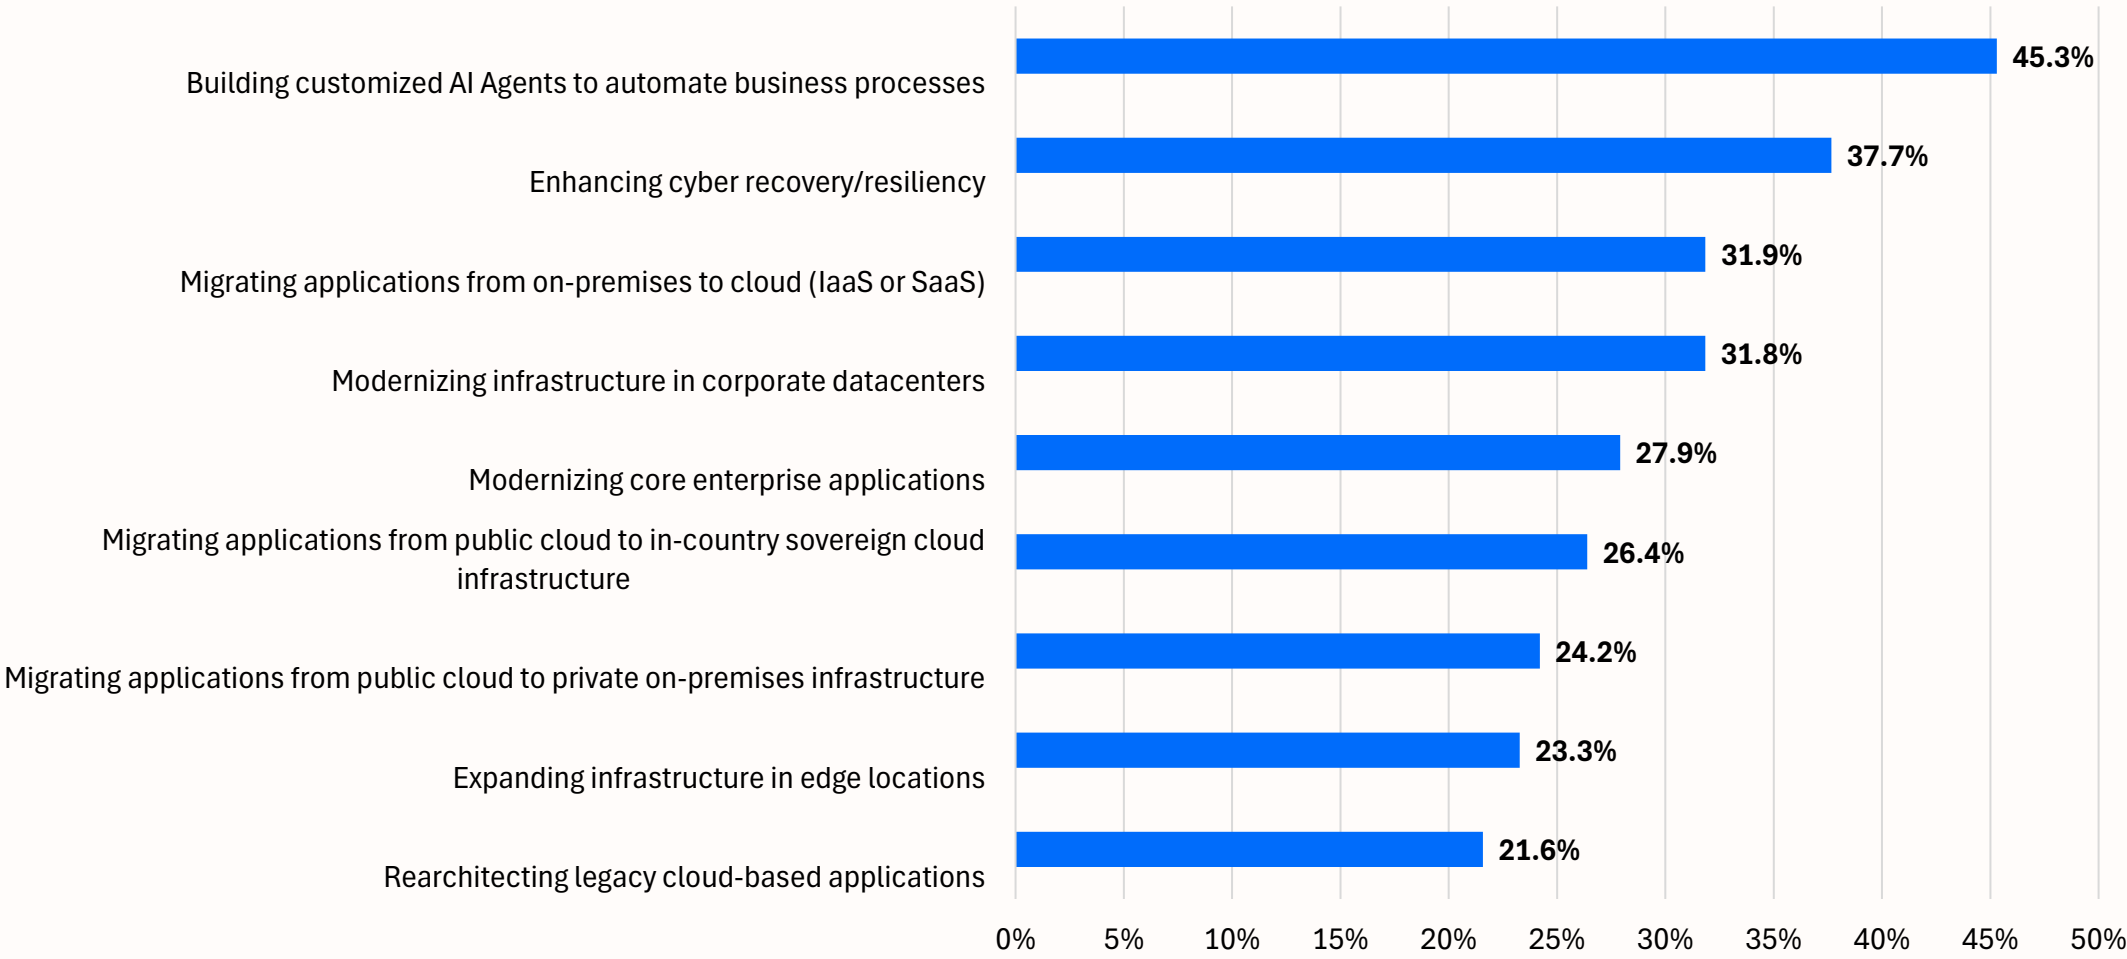


# Budgets are prepared for winning the gold. QITBQ4. Based on your current expectations for final 2025 spending, what is your best estimate of your organization's likely IT spending budget for 2026?




Source: IDC Syndicated Survey 2025: Future Enterprise Resiliency & Spending Wave 10, October, 2025. N=879.

# QITBQ10. Regardless of overall IT spending plans, in which of the following areas do you expect to most significantly increase spending in 2026?



Source: IDC Syndicated Survey 2025: Future Enterprise Resiliency & Spending Wave 9, September, 2025. N=1,007

**Don't let a recession ruin your cross-country run for gold.** QRVO1. A recession is defined as a period of temporary economic decline during which trade and industrial activity are reduced, generally identified by a fall in GDP for an extended period. Thinking about the economic situation in your country, how likely do you think it will be for a recession to happen in the next 12 months?




Source: IDC Syndicated Survey 2025: Future Enterprise Resiliency & Spending Wave 10, October, 2025. N=879.



# Market Forecast goes for Ice Skating **glory**

Worldwide Intelligent CloudOps Software Revenue by Functional Market, 2024–2029



**Total Market**  
2024–2029 CAGR  
14.0%

**Automation and configuration**  
2024–2029 CAGR  
13.0%

**Operations management**  
2024–2029 CAGR  
13.8%

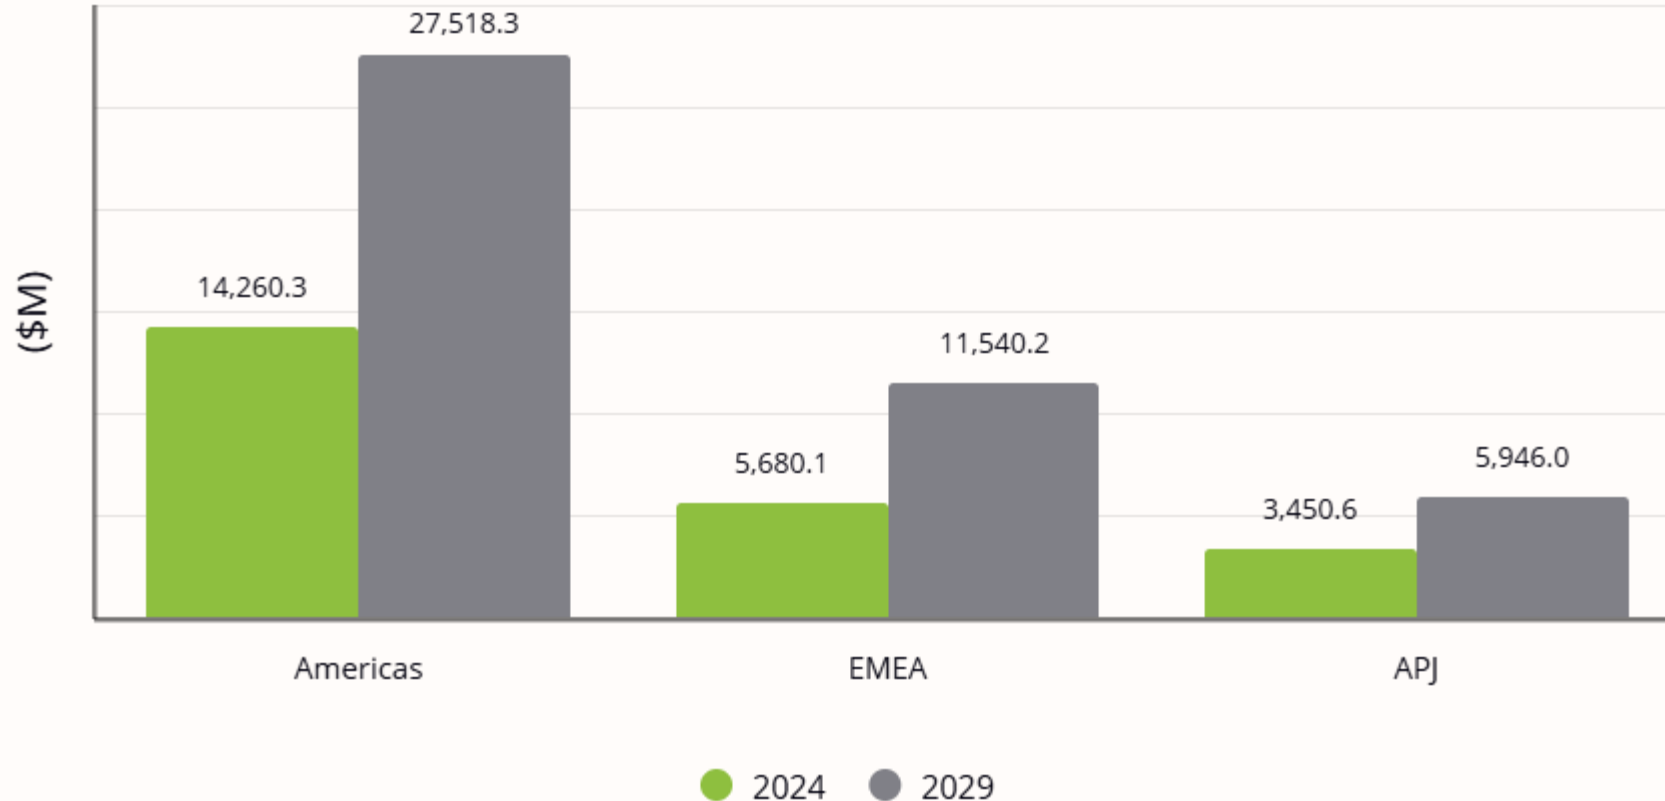
**Service management**  
2024–2029 CAGR  
16.1%

Source: IDC's Semiannual Software Tracker, 2H24



# Market Forecast

## Worldwide Intelligent CloudOps Software Revenue by Region, 2024 and 2029 (\$M)



Region	2024–2029 CAGR
Americas	14.1%
EMEA	15.2%
APJ	11.5%

Region	2024 Share
Americas	61.0%
EMEA	24.3%
APJ	14.8%


Region	2029 Share
Americas	61.1%
EMEA	25.6%
APJ	

Source: IDC's Semiannual Software Tracker, 2H24



# Executive Summary

## Worldwide Cloud Cost Transparency Software 2024 Share Snapshot




- Strong, steady growth has pushed the market to over \$2 billion for the first time.
- Acquisitions propel Flexera into the top 3.
- Growth slowed from the prior year slightly, but it still outpaced the larger software market. IDC further refined its model, increasing past-year revenue figures, which may impact year-over-year growth.

Source: IDC, November 2025




# Cost Pressures Drive Adoption and Investment in FinOps




**87.9%**

Percentage of Enterprises Utilizing FinOps



**8**

Median Size of FinOps Teams



**13**

Median Size of Cloud Center of Excellence Teams

Source: United States — IDC's Intelligent Cloud Operations Survey, November 2025 (n = 102)

Distribution is forbidden without written consent of IDC. All rights reserved.



# Why SaaS & PaaS Platforms Matter to FinOps



SaaS and PaaS now dominate IT spending



Data Platforms drive new growth



Concentrated Business Value needs to be in scope



Usage and Tiered in SaaS pricing creates risks



Cost telemetry methods vary by provider



All platforms need to be in scope for full cloud spend visibility



# Advice #1: *Mature FinOps* practices for business value.

Combining IT automation, AI, and open collaboration accelerates maturity



## In 2026, IT and business leaders should:

- **Standardization:** Benchmark cloud spending and ROI to other companies, not just internally to historical figures. Companies often settle for incremental improvements in a few metrics and maturity levels rather than stretch goals that make greater strides. Compare your FinOps maturity against peers and competitors.
- **Platform-centricity:** Take a platform approach to expand the FinOps discipline to the full stack and successfully optimize spend. Combine insights, recommendations and automation to implement recommendations. Today, the use of IT automation for FinOps recommendations is limited.
- **Leverage Budgets to invest in FinOps training and Tool upgrades:** Invest in people and FinOps Foundation certification to bring best practices back to your organization. Integrated with DevOps and other teams for improved transparency.
- **Expand Scope of Teams:** Include private cloud costs, SaaS, and Software asset management as ways to unlock new savings for the business. Share wins with executives for improved sponsorship and ensure company-wide accountability, responsibility, and transparency.



# Advice #1: Mature FinOps Today. Real World Example.



**Focus:** By Reaching 90%+ cost allocation accuracy and integrating FinOps into daily operations. Key metric to move to 'Run' Maturity.



**Business Value:** IT operates as a **strategic partner** rather than a cost center, using **unit economics** (e.g., cost per transaction) to prove ROI for new initiatives.



**Example:** A large healthcare provider matured their practices to right-size servers against specific business objectives, resulting in **\$1 million in annual savings** while maintaining operational efficiency



# AI ROI Impacted by New Cost Elements and Complexity

Need for **FinOps for AI** (projects)

“

*AI projects are at risk when the key drivers of its TCO are misunderstood.*

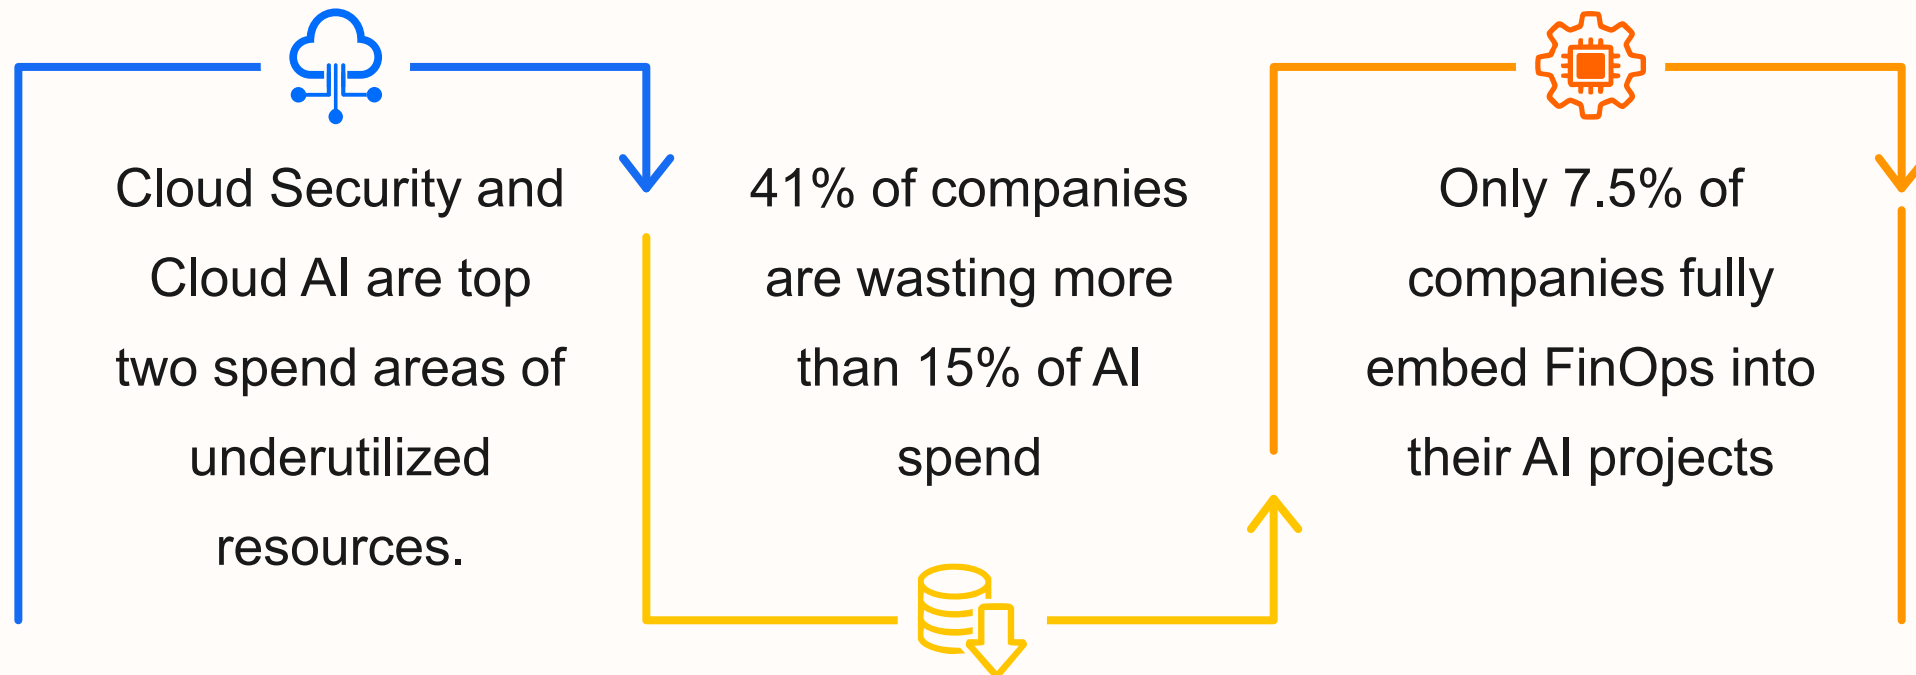
**Jevin Jensen,**  
**IDC Research VP**

- 1 New costs not found in traditional IT projects.
- 2 Infrastructure is not AI-ready, including high-speed networks and low-latency storage.
- 3 Proper workload placement consideration is required. On-premises is making a comeback with AI training and inference.
- 4 Legal reviews, proper governance, and new cybersecurity testing methods are utilized.
- 5 Cost allocation and chargeback of tokens or API usage increases complexity.
- 6 Master data management and ongoing model costs.



# FinOps for AI: Where to focus in 2026 on the win

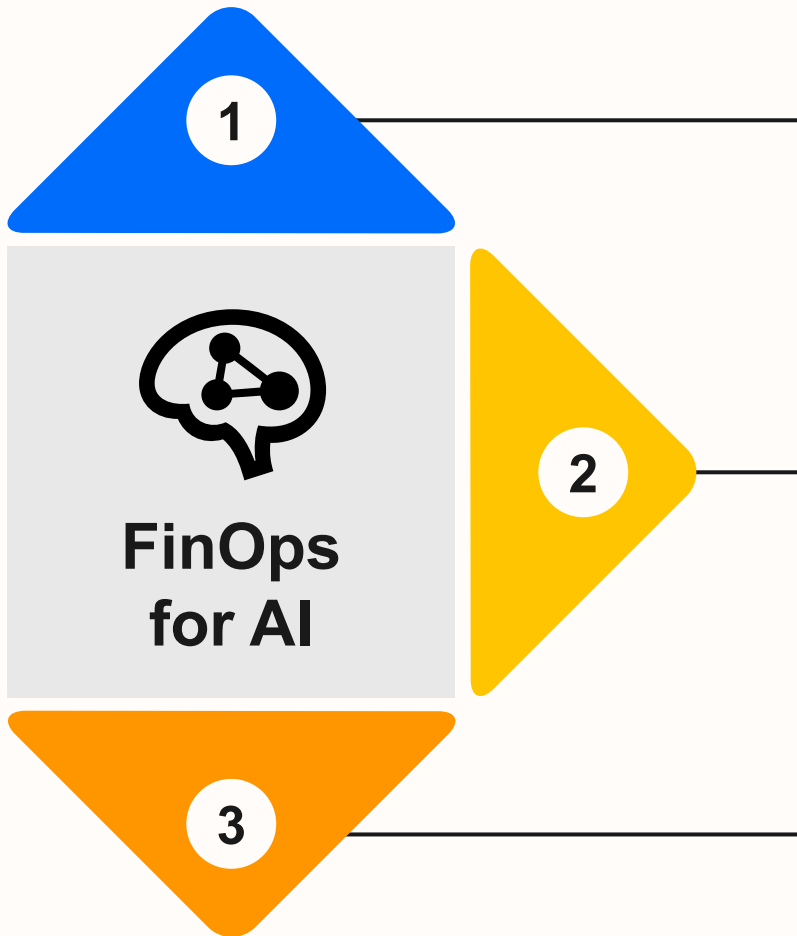
Build your own rings of success and mature your FinOps teams.



Source: IDC Europe, EMEA Cloud Ops Survey 2025, November 2025. N=509



# Advice #2: *FinOps for AI* for faster paybacks on projects



## AI differs in several essential ways – Educate your Teams!

- AI terms and concepts differ from typical cloud terms that FinOps teams are familiar with. Educate end users on the latest hyperscalers' AI tools/costs (Azure OAI, Google Cloud costs, Oracle Analytical Cloud, and AWS Cost Explorer for Bedrock).
- Continual training/optimization of models is necessary and means increased ongoing costs (e.g., pruning or distillation can reduce model optimization costs). AI agents tend to need more integration and orchestration costs to access internal/external model data.
- Engineering teams may be immature with proper AI and data hygiene and need additional coaching; new personas, such as data scientists, may need to be added to the FinOps team. New training and awareness programs for team members.

## Cloud cost optimization best practices for AI

- Push cloud providers for lower prices via savings plan recommendations for key AI elements and web services.
- Tagging remains key to accountability and proper cost allocation. AI resources can and should be tagged. Help teams automate tagging. Tag AI services like SageMaker, Bedrock, etc.
- Focus on cloud infrastructure savings like network, storage, and compute (GPU) optimization. These spikes are even higher than those in typical IaaS applications.

## Recommendations on tokens versus subscriptions

- Alert on high usage-based tokens versus time-based (100 reports of 500 words each = 50,000 tokens versus 2 CPU x 50 hours = \$5).
- Provide best practices on "input tokens" (prompts) versus generated tokens (responses).
- Add AI agents/copilots to implement cloud cost recommendations after approval from FinOps teams. Automate allocating AI API and AI agent usage



## Advice #2: FinOps for AI. Real World Example.



**Focus:** Multiple AI Agent pilots running in parallel. Used FinOps teams to review and prioritize each project. Applied optimization to GPU usage and associated sprints, while ensuring all costs were captured.



**Business Value:** High-ROI AI use cases remained fully funded while **cutting low-value experiments**. Improved the ROI of AI spend overall and freed up budget for new, revenue-generating AI features (e.g., faster claims handling and better risk pricing).




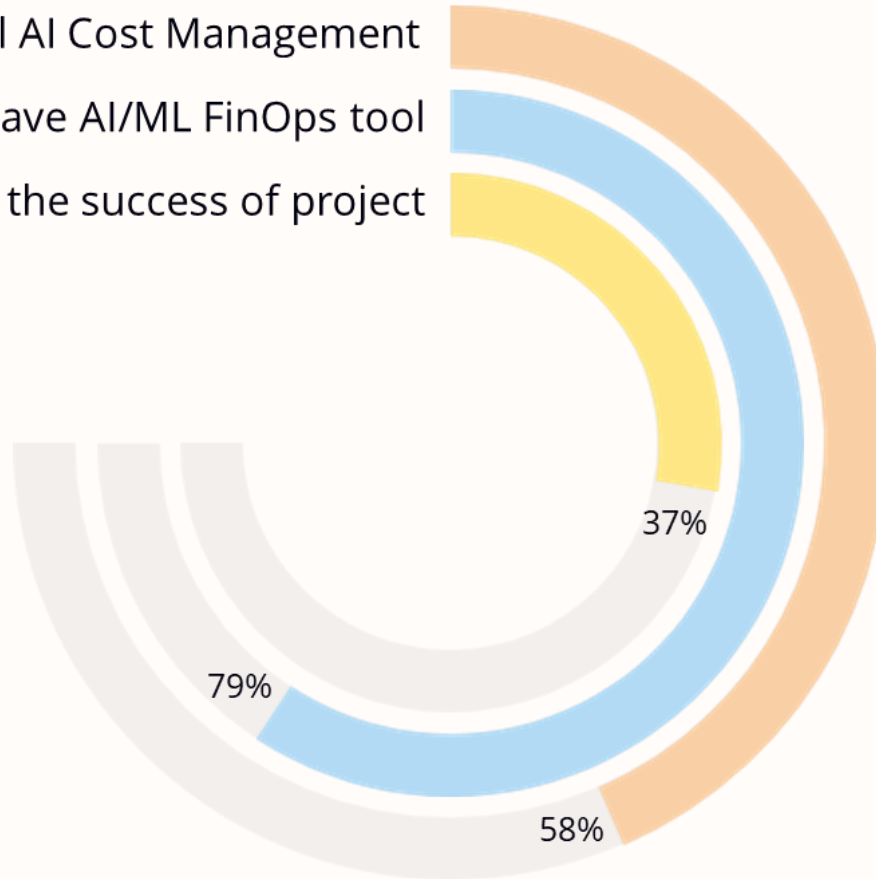
**Example:** A global insurer identified **about \$17 million in annual savings** opportunities, including **\$6 million** in “quick wins,” such as rightsizing GPU clusters and shutting down idle AI pilot sandboxes.



# Move forward with People AND Technology

Most important skillset in next 12 months and Importance of AI/ML in FinOps tool?

Top Ranked Skill AI Cost Management  
Very or Important to have AI/ML FinOps tool  
Lack AI skills impact the success of project




Source: Intelligent Cloud and Edge Operations with FinOps (Opportunity and Markets) Survey 2026, IDC, November 2026, N=102. FinOps Foundation, February 2026 Survey, N=695. IDC SAS: AI Trust Survey 2025, June, 2025. N=2,375.




# IDC Survey Spotlight

## Are FinOps teams using tools for AI-related cloud cost optimization?

Q8. Are you willing to pay more for an AI or GenAI feature in your current FinOps solution, do you use your FinOps tool to help optimize AI-related cloud costs?



The current tool supports optimizing AI-related costs.




Willing to pay for AI features in tool.

Usage-based

Hybrid

Subscription

License-based




Top 3 preferences to measure for AI-related charges?

Source: US – Intelligent Cloud Operations Survey 2025, IDC, October 2025. N = 107




# Where are companies growing data and production workloads?

Q38. Please estimate the annual data growth rate for each of the following environments within your organization over the past 12 months.




Number of companies expected to grow PUBLIC cloud by more than 10%



Number of companies expected to grow PRIVATE cloud by more than 10%

What is the approximate percentage of your organization's total IT PRODUCTION transactions processed by location type today?




IDC #US54233326 (January 2026)  
Source: IDC's Europe, EMEA CloudOps Survey 2025, November 2025 (n = 509) and IDC's U.S. Intelligent Cloud Operations Survey 2025, September 2025 (n = 107)  
Distribution is forbidden without written consent of IDC. All rights reserved.

## Advice #3: Focus on *AI for FinOps* to control costs.

- Applying the recommendations from a FinOps tool faster means greater savings to the bottom line. Using AI with automation creates 'smart automation' to get those implementations quicker and safer with guardrails.
- Forecasting and what-if scenarios are great uses of AI found in some FinOps platforms. Using historical and real-time detection to predict the future while alerting on spend anomalies before getting your hyperscaler bill saves money.
- Breaks down silos for business leaders to more easily utilize the FinOps platform. Utilizing GenAI to ask natural language questions about spending, create reports, or update dashboards increases transparency

**Q. Top two strategies for controlling costs?**



**Answer:** Automating AI portfolio monitoring, life-cycle management, and reporting.

Source: IDC's *Future Enterprise Resiliency and Spending (FERS) Survey, Wave 6, June, 2025*. (n = 920)



## Advice #3: AI for FinOps. Real World Example.



**Focus:** Deployed an AI-enabled FinOps strategy that combined systematic cloud optimization with automated right-sizing and EBS volume migration.



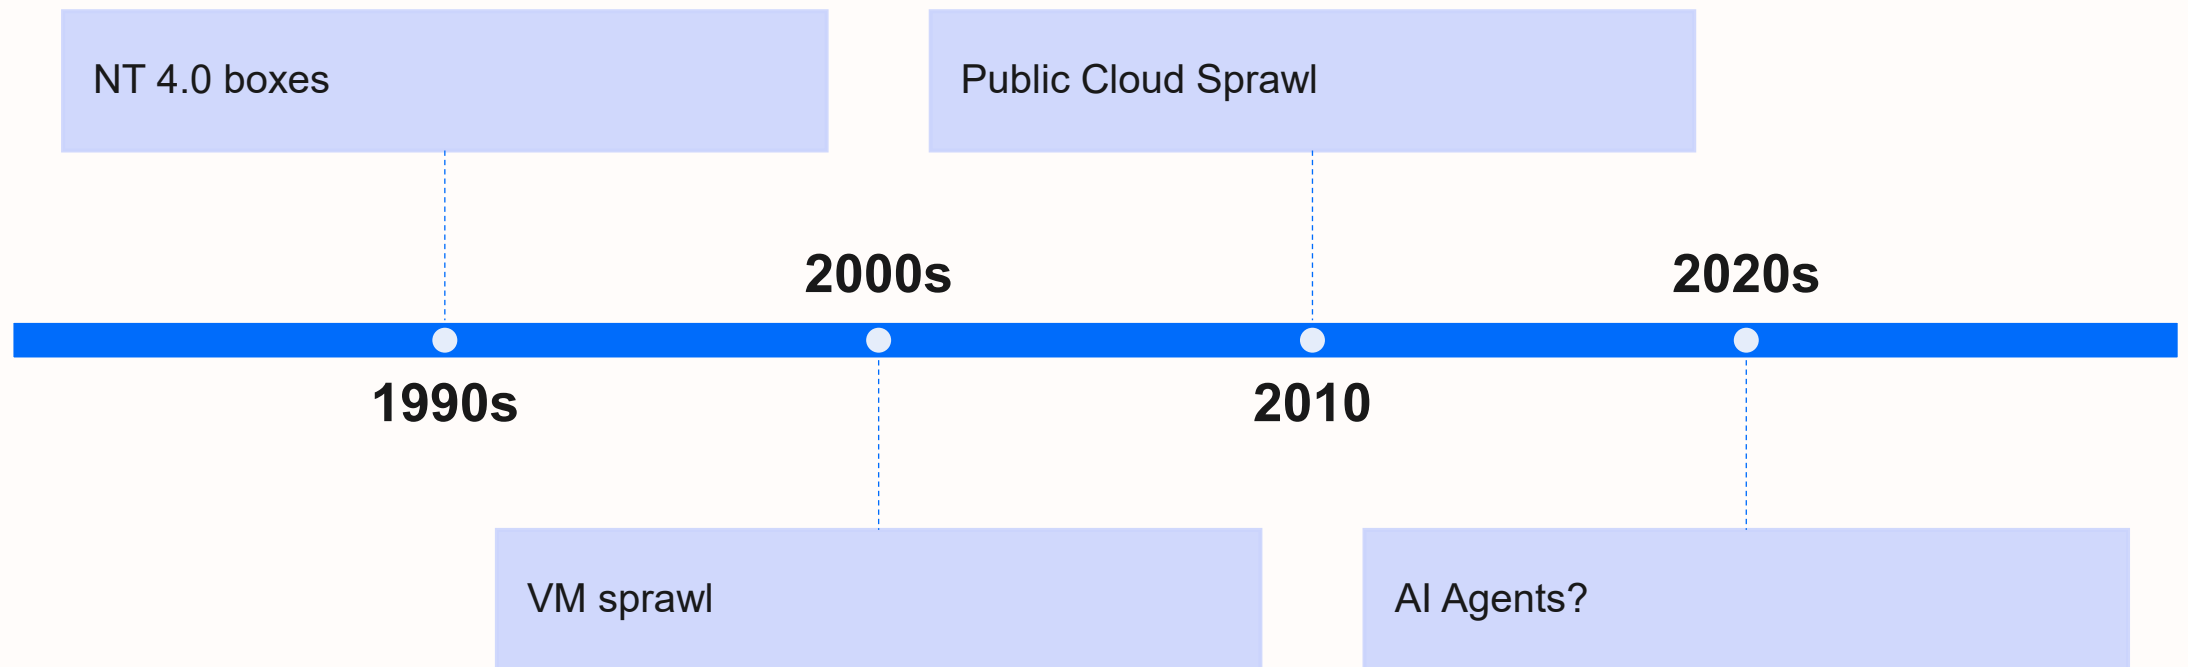
**Business Value:** FinOps tools that offer AI/ML insights can provide better, faster recommendations. Adding agents or other real-time automation to the process can resolve issues faster than monthly report-based recommendations.



**Example:** ERP software firm realized over **\$2.3 million in savings** by eliminating overprovisioning and improving storage efficiency. EC2 savings



# IT has a checkered history with new technology adoption




# The era of agent economics



## How many agents

(28.8 million agents in 2025)

- 80x more agents than 2025
- Agents double again in 2030
- 35% are customizable agents (low code/no code)
- 2.3 Billion Agents by 2030
- 10,000 employee enterprise would have 6,700 agents



## How many actions

(120 million actions per day in 2025)

- Agent actions per day will approach **217 billion** in 2029
- 1,798x more actions
  - Actions increase 4x in 2030
  - 33% are by customizable agents



## Will you be ready?

- Agent2Agent Communication? MCP?
- Shadow IT with AI Agents bought from marketplace?
- Inventory of all agents, how to monitor and maintain resilience?
- Can business continue if an Agent fails? DR for Agents?



## Advice #4: AI Agent Sprawl. Real World Example.



**Focus:** Estimated **1.5 million unmonitored AI Agents** today. Forecast to get much worse as ease of procuring (shadow IT is back!) or creating new Agents requires limited technical knowledge.



**Business Value:** AI agents may become **embedded in mission-critical** processes overnight. Niche agents may double or triple an enterprise's agent usage overnight. Cybersecurity and resiliency issues with unmanaged agents will multiply.

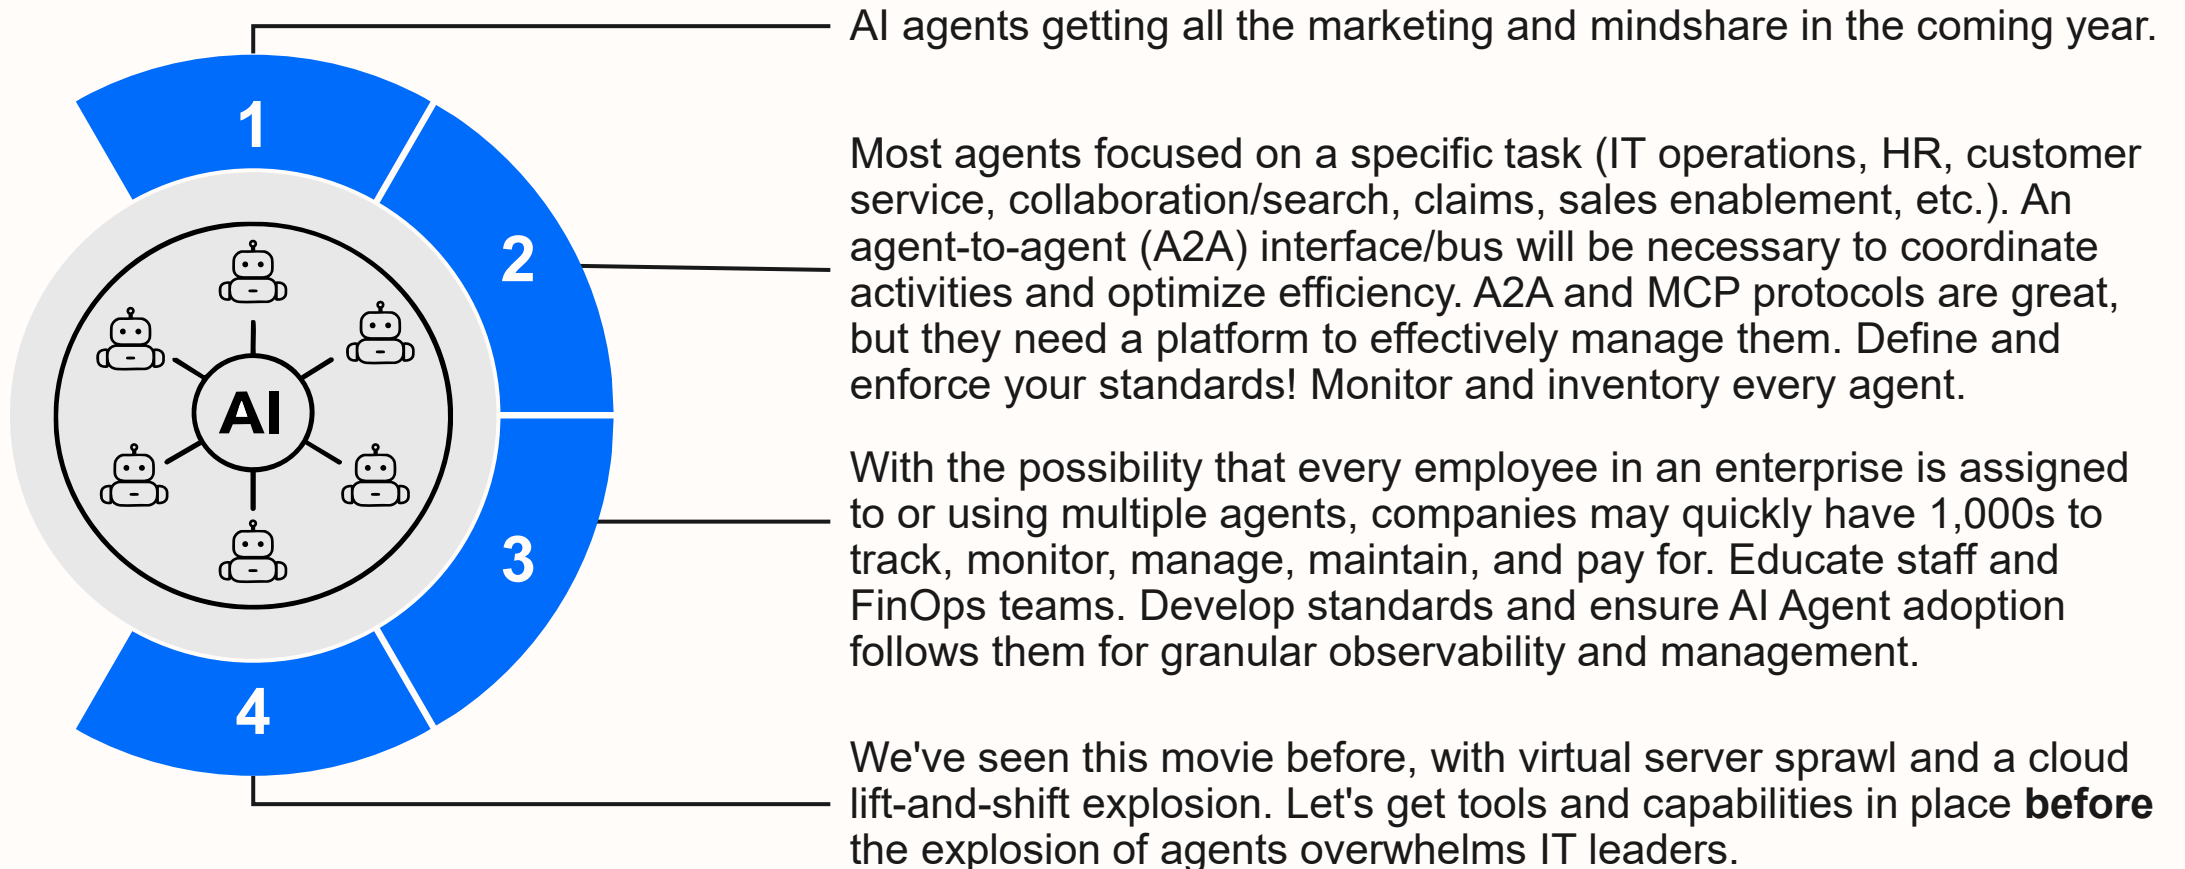


**Example:** Don't repeat the mistakes of the past. Unmonitored agents that cause outages and support issues risk losing all efficiency gains from their usage.



# Advice #4: *AI Agent Sprawl* is coming soon...

The movie no one wants to see again! History repeats itself.



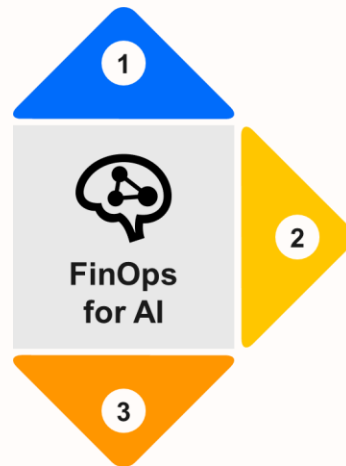
# Conclusion

Four takeaways for your business.

Mature FinOps Teams



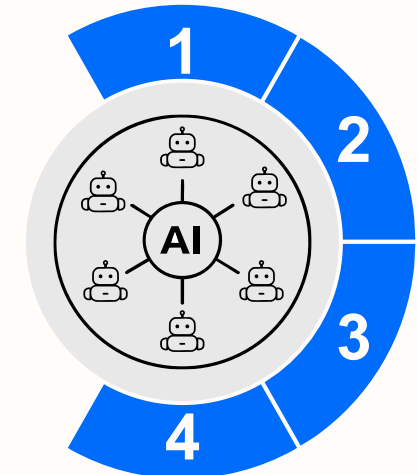
FinOps for AI



AI for FinOps



Prepare for AI Agents



# Jeivin Jensen

[jjensen@idc.com](mailto:jjensen@idc.com)

[linkedin.com/in/jevinjensen](https://www.linkedin.com/in/jevinjensen)

