


Gartner Insights

A CIO's Guide to AI Commercial and Pricing Models

Shruthi Shankel, Hannah Decker, Bill Ryan,
John Santoro, Lauren Wheatley, Akanksha Pandey,
Andy Rowsell-Jones



A CIO's Guide to AI Commercial and Pricing Models

5 May 2026 - ID G00851577 - 12 min read

By: Shruthi Shankel, Hannah Decker, Bill Ryan, John Santoro, Lauren Wheatley, Akanksha Pandey, Andy Rowsell-Jones

Initiatives: Financial; Manage Costs Strategically to Spend Smarter and Scale Faster; Accelerating Enterprise AI Value Realization at Scale

AI pricing models are evolving rapidly and vary widely across vendors. This research helps CIOs interpret and compare AI commercial and pricing models by defining common structures and terminology, creating a foundation for informed evaluation and future financial governance decisions.

Insights at a Glance

CIOs are used to comparing licenses, usage tiers, and enterprise agreements; however, AI is reshaping the market as vendors offer an already-complex and growing mix of commercial and pricing models, often blended together in hybrid forms.

CIOs Face a New Reality

AI pricing proposals are opaque, inconsistent, and increasingly difficult to compare. The result is uncertainty around cost predictability, value realization, and scalability. This is no longer a procurement issue but a financial governance one. Without a consistent way to interpret AI pricing, enterprises risk losing cost predictability, misaligning spend and value, and misunderstanding the true economic impact of AI adoption.

To establish clarity in a fragmented and evolving AI pricing landscape, CIOs must define a common language for how AI is commercialized and priced. This starts with two foundational steps:

1. **Identify the commercial model:** the type of commitment being proposed (subscription, committed consumption, or pay as you go).
2. **Identify the pricing model:** how cost is calculated, including access-based, usage-based, activity-based, and outcome-based pricing, often used in combination.

Understanding AI pricing is complex and requires clarity on how commercial commitments, pricing models, and pricing metrics interact to shape cost behavior at scale. This is explored in [Breaking Down AI Pricing Complexity](#).

Strategic Planning Assumption

By 2028, 50% of infrastructure and operations (I&O) leaders will require AI agent performance SLAs in vendor contracts, up from less than 5% in 2025 — signaling a shift toward outcome-based pricing.

Issue

CIOs face a new reality: AI pricing proposals are opaque, inconsistent, and increasingly difficult to compare.

Enterprise AI adoption is accelerating rapidly. In the 2026 Gartner CIO and Technology Executive Survey, 52% of CIOs and technology executives reported that their organization had already deployed AI, with the remainder expecting deployment within the next two to three years. ¹ At the same time, vendors are introducing an expanding range of AI pricing approaches — from access-based pricing embedded in enterprise software to token-based model consumption, activity-based automation, and emerging outcome-based services.

This complexity is compounded by low enterprise confidence in core financial disciplines related to AI. Recent Gartner webinar polling shows that the majority of CIOs and senior leaders report low or no confidence in their ability to measure AI ROI (63%), estimate AI costs (45%), or identify and assess AI-related risks (42%). ² As AI spend shifts from predictable licenses to dynamic, consumption- and value-based models, these confidence gaps become increasingly consequential.

As AI capabilities expand across platforms, applications, and autonomous agents, pricing models lack standard units of measure and consistent commercial constructs. Unlike traditional enterprise software, AI pricing makes it harder for enterprises to interpret contracts, compare offers across vendors, and anticipate how costs will scale across use cases.

CIOs must therefore evaluate not only the price of AI services, but also how different pricing and commercial models shape cost behavior, scalability, and financial accountability as adoption grows. Without a common language for AI pricing models, enterprises struggle to make coherent decisions, govern spend consistently, and sustain confidence in AI investments.

Impact

AI pricing complexity directly affects how enterprises plan, justify, and govern AI spending. When commercial commitments and pricing structures are poorly understood, organizations struggle to answer basic questions about what drives AI cost, how spending scales across use cases, and why costs vary among vendors.

As AI adoption expands, this ambiguity disrupts financial planning and accountability. Budgeting and forecasting processes built for subscription-based access models or predictable unit pricing are strained by pricing tied to tokens, activities, or outcomes that scale unevenly across teams and workflows. This reduces confidence in cost estimates, complicates chargeback and allocation, and makes AI spend harder to justify to finance and executive stakeholders.

Pricing opacity also weakens vendor accountability. When cost drivers are unclear or embedded across multiple pricing constructs, enterprises have limited ability to benchmark offers, compare alternatives, or link spend to delivered value. Over time, this increases financial exposure and leaves organizations with AI cost structures they cannot clearly explain or confidently govern as AI adoption accelerates.

Understanding AI Commercial and Pricing Models

Understanding AI pricing starts with two steps: the commercial model that defines commitment and the pricing model that drives cost. Without separating these, AI pricing remains difficult to interpret or compare.

Step 1: Identify the Commercial Model – What Are You Committing To?

There are a variety of commercial models that vendors use to sell their products. The first step in understanding what you have, or what you need, is to determine the commercial model. There are three options: traditional subscription, committed consumption and pay as you go.

Traditional Subscription

This is the most familiar commercial structure:

- It involves a monetary commitment based on a defined number of users, seats, or licenses in exchange for access to specified products for a fixed period.
- Pricing is determined by the number of users or licenses multiplied by a recurring (typically annual) fee.
- Billing is usually an annual upfront payment or fixed periodic payments (monthly or quarterly).

Examples: Microsoft Copilot (per user/month), Google Workspace add-ons (per user or per seat/license per month), and Zapier (which is sold as a subscription with pricing tiers that reflect usage limits). ^{3,4,5}

Committed Consumption

This is a commitment to spend rather than to specific products or licenses:

- It involves a commitment to a defined level of spend or credits (such as tokens, APIs, or compute) that can be drawn down over a fixed period.
- Commitments may apply to a **specific product** or to a **portfolio of services** under a broader agreement.
- Unused consumption is typically forfeited at the end of the term unless contractually negotiated.
- Overconsumption beyond the committed amount is permitted but charged separately, often at on-demand rates or with reduced discounts.
- Billing is typically an annual upfront commitment or monthly billing against a committed pool, with end-of-term reconciliation (true-up for overusage or recognition of unused consumption).

Examples: Microsoft Azure consumption commitments, OpenAI enterprise credit-based plans, AWS Enterprise Discount Program (EDP), SAP Business Technology Platform Enterprise Agreement (BTP EA). ^{6,7,8,9}

Pay As You Go

This model involves no upfront commitment; customers are charged only for what they consume:

- Customers have no commitment to a minimum spend, credits, or product quantities.
- Costs accrue based on actual usage of services, APIs, or tokens.
- Billing is typically in arrears, based on measured consumption during the billing period.

Examples: Google Vertex AI usage-based billing, OpenAI API usage (per token), Anthropic Claude API usage. ^{10,11,12}

Step 2: Identify the Pricing Model – What Drives Cost?

Once commitment is clear, the real challenge begins: What are you actually paying for? In reality, most AI offerings rely on four foundational pricing models – access-based, usage-based, activity-based and outcome-based. CIOs should start by identifying the underlying pricing structure, as it directly determines how costs scale, who owns risk, and how easily spending can be controlled as adoption grows. Without this clarity, AI cost behavior becomes difficult to anticipate and explain as adoption expands.

Table 1 outlines the four AI pricing models, showing how enterprises are charged by the vendors, the units used to determine cost (pricing metrics), and the implications for the enterprise. While usage- and activity-based models are often conflated, usage pricing charges for AI resource consumption, whereas activity-based pricing charges for completed work or transactions, a distinction that materially affects cost behavior and risk allocation. Examples provided in the metrics are purely illustrative. Be aware that vendors may apply these metrics in different combinations. Together, these elements help CIOs understand cost behavior, risk, and trade-offs across different pricing approaches.

Table 1: Four Foundational AI Pricing Models

(Enlarged table in Appendix)

Pricing Model	Enterprise Cost and Risk Implications	Example Pricing Metrics
Access-Based A fixed recurring fee for user- or platform-level access to AI capabilities, independent of actual usage.	- Predictable cost; however, underutilization risk sits with the buyer. - Encourages broad access but not necessarily efficient usage. - Can lead to tool sprawl and overlapping capabilities across vendors.	- Per user/per seat, e.g., per workspace - AI add-on fee (e.g., \$15 per user, per month added to an existing SaaS license for new AI features) - Tier/Edition/Plan-based feature unlock (e.g., Basic/Pro/Enterprise AI)
Usage-Based Charges based on how much AI infrastructure or model capacity is used.	- Provides flexibility, but budget volatility risk sits with the buyer as it scales with usage. - Risk of hidden cost drivers (e.g., inefficient prompts, repeated calls, background processing). - Requires real-time visibility and cost control mechanisms as adoption grows.	- Tokens - API calls - Compute hours - Credits - Storage used - Automation runs
Activity-Based Charges are tied to tasks or units of work completed by AI systems. This is an evolving model with adoption rising across vendors.	- Costs scale with operational activity rather than raw infrastructure usage. - Predictability depends on transaction volumes. - Vendors assume execution-efficiency risk, while enterprises retain demand and volume risk. - Requires clear process definitions and strong controls over activity levels.	- Per invoice/document processed - Per ticket handled - Per video/image generated - Per contract reviewed - Per chatbot interaction handled - Per claim processed - Per workflow executed
Outcome-Based Charges are based on measurable business outcomes delivered by AI systems, not usage or activity alone. While adoption remains limited, outcome-based pricing is emerging as vendors experiment with performance- and value-linked models.	- This is not just work completed, but also value delivered against explicitly defined and measurable outcomes. - Vendors assume performance risk, which is typically reflected in pricing and contract structure. - Success depends on high-quality data, clear baselines, and robust attribution models. - Risk of disputes increases where outcomes are ambiguous or materially influenced by external factors. - Best-suited to mature use cases where outcomes are well-defined and directly attributable.	- Percent revenue recovered (e.g., 10% of unpaid invoices) - Percent fraud prevented (e.g., 15% of blocked fraudulent transaction value) - Percent of cost saved - Percent of conversion uplift - Commission on sales generated (e.g., 5% AI-driven sales) - Gainshare-/performance-based contracts (e.g., pay tied to KPIs: 95% accuracy, 30% faster resolution time)

Source: Gartner (May 2026)

In practice, enterprise AI pricing often combines multiple foundational models. For example, an AI platform may charge a base subscription price for platform access while also charging consumption fees for model usage or activity-based fees for automation workflows. Four current and emerging hybrid AI pricing model structures are outlined in Table 2.

Table 2: Common Hybrid AI Pricing Models

(Enlarged table in Appendix)

Hybrid Pricing Model	Implications for the Enterprise	Example Pricing Metrics
Access + Usage Base platform subscription combined with usage-based charges.	- Costs increase as AI workloads scale across the enterprise. - Enterprise carries most of the financial risk, particularly if AI adoption grows rapidly or usage patterns are difficult to predict.	- Per user/seat + vendor credits - Per enterprise license + token usage - Per user/seat + API calls - Enterprise license + compute credits (e.g., \$200,000 annual enterprise license + \$2 per compute hour or \$1 per GB processed)
Access + Activity Platform access fee combined with pricing tied to tasks executed by AI systems.	- Risk begins to shift slightly toward the vendor because costs are linked to operational output rather than raw infrastructure usage. - Enterprises gain better alignment between AI activity and spend.	- Per user/seat + per workflow run - Per automation license + document processed (e.g., \$30,000/year automation license + \$0.10 per document processed for extraction)
Usage + Activity Compute usage combined with task-based pricing.	- Costs scale based on both compute usage and tasks executed. - Risk is balanced between buyer and vendor because pricing reflects both infrastructure demand and actual AI work performed. - Enterprises must monitor usage closely to prevent cost escalation as automation expands.	Per API calls/tokens + per document/transaction/workflow processed (e.g., \$0.0003 per 1000 tokens processed + \$0.50 per document extracted or \$0.02 per chatbot interaction)
Activity + Outcome Charge based on tasks completed plus incentives tied to outcomes.	- Vendors carry greater performance risk because part of the payment depends on the results delivered. - For enterprises, this model creates stronger alignment between AI cost and business value, but requires clear governance and outcome measurement.	- Per case processed + revenue share (e.g., \$5 per case processed + 5% of additional revenue generated from those cases) - Per lead generated + conversion-based commission (e.g., \$5 per lead generated + 4% of commission on sales from leads that successfully convert)

Source: Gartner (May 2026)

AI Pricing Model Landscape: Current and Emerging Models

These distinctions explain why AI pricing often appears fragmented despite similar underlying capabilities. Vendors apply fundamentally different pricing logics to comparable use cases, making headline price comparisons misleading.

The AI pricing landscape spans access-based, usage-based, activity-based and outcome-based models, frequently combined in hybrid structures. In practice, most enterprise AI offerings combine multiple pricing models rather than relying on a single construct. Figure 1 illustrates sample product names, with vendors in brackets. The figure is illustrative, not exhaustive, and is intended to clarify market structure rather than represent a complete vendor set.


Sample products and vendors include:

- ChatGPT Plus (OpenAI) for access-based pricing
- Bland AI (Bland) for usage-based pricing
- Decagon AI (Decagon) for activity-based pricing
- AirHelp AI (AirHelp) for outcome-based pricing
- Adobe Creative Cloud (Adobe) for hybrid pricing

Product and vendor classification follows the pricing definitions in Tables 1 and 2: access-based (per user or platform access), usage-based (per token, API call or credits), activity-based (per task or workflow completed), outcome-based (share of uplift or cost saved), and hybrid models that blend multiple approaches.

Figure 1: AI Pricing Model Landscape

AI Pricing Model Landscape



Vendor offerings are segmented based on the dominant AI pricing models observed in publicly disclosed offerings. Many vendors support multiple models. Placement reflects Gartner analysis at the time of publication, is illustrative rather than exhaustive, and does not imply evaluation or comparative judgment.

Source: Gartner
851577

Gartner

Gartner 2026

As enterprises scale AI adoption, the limits of traditional cost management become clear. Variable consumption, shared services, and outcome-based models strain budgeting, forecasting, and value tracking. This shift exposes a deeper governance gap, one that becomes most visible in how AI is priced.

Contributors

Chris Ganly

Evidence

¹ 2026 Gartner CIO and Technology Executive Survey. This survey was conducted online from 1 May to 30 June 2025 to help CIOs and technology executives benchmark their priorities and investment plans against those of peers worldwide. Qualified respondents led a digital/technology function and were accountable for running or improving/growing a specific area of their enterprise. In total, 2,501 CIOs and technology executives participated, with representation from all geographies, revenue bands and industry sectors (public and private).

² Polling data from AI Value Creation: Achieving ROI Demands a Labor-based Mindset Shift webinar, 19 June 2025

³ Microsoft 365 Copilot Plans and Pricing — AI for Business, Microsoft

⁴ Understand Google Workspace bills & charges | Billing & subscriptions, Google Workspace

⁵ AI Orchestration Plans, Zapier

⁶ Azure Pricing Overview, Microsoft

⁷ Flexible Pricing for the Enterprise, Edu, and Business Plans, OpenAI Help Center

⁸ AWS EDP: How It Works and Why Automation Will Give You Better Results, cast.ai

⁹ Pricing | SAP Business Technology Platform, SAP

¹⁰ Vertex AI consumption options | Generative AI on Vertex AI | Google Cloud Documentation, Google cloud

¹¹ OpenAI API Pricing

¹² Pricing - Claude API Docs

Recommended by the Authors

Some documents may not be available as part of your current Gartner subscription.

[Maximize Value From GenAI Services With Optimal Commercial Model Alignment](#)

[Breaking Down AI Pricing Complexity](#)

[Buyer Beware: Evaluating AI Pricing in Vendor Selection](#)

[Prevent AI Budget Shock With Pricing Governance](#)

[Pricing Models Based on AI Units Are Essential to Adoption and Profits](#)

[AI Disruption & SaaS Slowdown: The Return of Buyer Leverage](#)

© 2026 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's Business and Technology Insights Organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner insights may address legal and financial issues, Gartner does not provide legal or investment advice and its insights should not be construed or used as such. Your access and use of this publication are governed by [Gartner's Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its insights is produced independently by its Business and Technology Insights Organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)." Gartner insights may not be used as input into or for the training or development of generative artificial intelligence, machine learning, algorithms, software, or related technologies.

Table 1: Four Foundational AI Pricing Models

Pricing Model	Enterprise Cost and Risk Implications	Example Pricing Metrics
Access-Based A fixed recurring fee for user- or platform-level access to AI capabilities, independent of actual usage.	- Predictable cost; however, underutilization risk sits with the buyer. - Encourages broad access but not necessarily efficient usage. - Can lead to tool sprawl and overlapping capabilities across vendors.	- Per user/per seat, e.g., per workspace - AI add-on fee (e.g., \$15 per user, per month added to an existing SaaS license for new AI features) - Tier/Edition/Plan-based feature unlock (e.g., Basic/Pro/Enterprise AI)
Usage-Based Charges based on how much AI infrastructure or model capacity is used.	- Provides flexibility, but budget volatility risk sits with the buyer as it scales with usage. - Risk of hidden cost drivers (e.g., inefficient prompts, repeated calls, background processing). - Requires real-time visibility and cost control mechanisms as adoption grows.	- Tokens - API calls - Compute hours - Credits - Storage used - Automation runs
Activity-Based Charges are tied to tasks or units of work completed by AI systems. This is an evolving model with adoption rising across vendors.	- Costs scale with operational activity rather than raw infrastructure usage. - Predictability depends on transaction volumes.	- Per invoice/document processed - Per ticket handled - Per video/image generated

- Vendors assume execution-efficiency risk, while enterprises retain demand and volume risk.
- Requires clear process definitions and strong controls over activity levels.
- Per contract reviewed
- Per chatbot interaction handled
- Per claim processed
- Per workflow executed

Outcome-Based

Charges are based on measurable business outcomes delivered by AI systems, not usage or activity alone. While adoption remains limited, outcome-based pricing is emerging as vendors experiment with performance- and value-linked models.

- This is not just work completed, but also value delivered against explicitly defined and measurable outcomes.
- Vendors assume performance risk, which is typically reflected in pricing and contract structure.
- Success depends on high-quality data, clear baselines, and robust attribution models.
- Risk of disputes increases where outcomes are ambiguous or materially influenced by external factors.
- Best-suited to mature use cases where outcomes are well-defined and directly attributable.
- Percent revenue recovered (e.g., 10% of unpaid invoices)
- Percent fraud prevented (e.g., 15% of blocked fraudulent transaction value)
- Percent of cost saved
- Percent of conversion uplift
- Commission on sales generated (e.g., 5% AI-driven sales)
- Gainshare-/performance-based contracts (e.g., pay tied to KPIs: 95% accuracy, 30% faster resolution time)

Source: Gartner (May 2026)

Table 2: Common Hybrid AI Pricing Models

Hybrid Pricing Model	Implications for the Enterprise	Example Pricing Metrics
Access + Usage Base platform subscription combined with usage-based charges.	■ Costs increase as AI workloads scale across the enterprise. ■ Enterprise carries most of the financial risk, particularly if AI adoption grows rapidly or usage patterns are difficult to predict.	■ Per user/seat + vendor credits ■ Per enterprise license + token usage ■ Per user/seat + API calls ■ Enterprise license + compute credits (e.g., \$200,000 annual enterprise license + \$2 per compute hour or \$1 per GB processed)
Access + Activity Platform access fee combined with pricing tied to tasks executed by AI systems.	■ Risk begins to shift slightly toward the vendor because costs are linked to operational output rather than raw infrastructure usage. ■ Enterprises gain better alignment between AI activity and spend.	■ Per user/seat + per workflow run ■ Per automation license + document processed (e.g., \$30,000/year automation license + \$0.10 per document processed for extraction)
Usage + Activity Compute usage combined with task-based pricing.	■ Costs scale based on both compute usage and tasks executed. ■ Risk is balanced between buyer and vendor because pricing reflects both infrastructure demand and actual AI work performed.	■ Per API calls/tokens + per document/transaction/workflow processed (e.g., \$0.0003 per 1000 tokens processed + \$0.50 per document extracted or \$0.02 per chatbot interaction)

- Enterprises must monitor usage closely to prevent cost escalation as automation expands.

Activity + Outcome

Charge based on tasks completed plus incentives tied to outcomes.

- Vendors carry greater performance risk because part of the payment depends on the results delivered.
- For enterprises, this model creates stronger alignment between AI cost and business value, but requires clear governance and outcome measurement.
- Per case processed + revenue share (e.g., \$5 per case processed + 5% of additional revenue generated from those cases)
- Per lead generated + conversion-based commission (e.g., \$5 per lead generated + 4% of commission on sales from leads that successfully convert)

Source: Gartner (May 2026)

Actionable, objective insights

Position your organization for success. Explore these additional complimentary resources and tools on for CIOs:

eBook

Build an Agile AI Sourcing and Contracting Strategy

Build a smarter approach to AI procurement and vendor selection.

[Download Now](#)



Hub Page

Gartner AI Hub

Discover practical AI insights for confident decision making.

[Learn More](#)



Virtual Briefing

IT 2030: Empowerment or Chaos? CIO's Role in AI-Powered Democratization

Learn how CIOs can enable AI innovation while managing risk.

[Watch Now](#)



Tool

Gartner AI Maturity Assessment

Assess your AI maturity and build a roadmap to scale impact.

[Explore Tool](#)



Already a client? Get access to even more resources in your client portal.
[Log In ↗](#)

Connect with us

Get actionable, objective business and technology insights that drive smarter decisions and stronger performance on your mission-critical priorities.

Contact us to become a client:

U.S.: 1 855 322 5484

International: +44 (0) 3300 296 946

[Become a Client](#)

Learn more about Gartner for CIOs
gartner.com/chief-information-officer

Stay connected to the latest insights



Attend a Gartner conference
[View Conference](#)