

Enterprise AI Monetization: 2025 Scorecard



Executive Summary

For the last two years, the industry competed on feature announcements and roadmaps. In 2025, the basis of competition has shifted. Our analysis shows the market has split into two distinct strategic realities: The **Monetizers**, who have successfully re-architected their business models to capture AI value, and The **Bundlers**, who are deploying AI primarily to bundle value and prevent customers from leaving.

Four key competitive shifts define the 2025 landscape:

1. Outcome-Based Pricing is Winning the Wallet

The battle for budget is no longer about productivity promises but proven value capture. While many competitors are stuck in pilot purgatory (95% failure rate), ServiceNow offers a leading example of successful AI pricing. By effectively selling an est. ~60% premium for specific, automated workflows, they have demonstrated that customers will pay for outcomes. This success validates that the next dominant GTM motion is hybrid outcome-based monetization, leaving competitors who rely on flat, seat-based add-ons vulnerable to commoditization.

2. Vertical Integration Beats Software-Only Plays

The software-only advantage is eroding as the value chain compresses. Oracle offers the definitive case study in successful strategic pivoting. By vertically integrating infrastructure (OCI) with applications, they have captured the demand that pure-play SaaS vendors missed. With AI GPU consumption spiking over 300% and OCI revenue growing 52%, Oracle proves that owning the underlying compute provides a competitive moat that application-layer incumbents cannot easily replicate.

3. The "Per-Seat" Business Model is Under Threat

Product and Sales leaders face an existential GTM crisis: AI agents are designed to reduce headcount, which cannibalizes the traditional per-seat revenue model. Salesforce and Microsoft are currently wrestling with this friction. Salesforce's aggressive pivot to consumption credits is not just a pricing tweak; it is a signal that the entire industry must prepare to transition from monetizing access (seats) to monetizing work (consumption/actions).

4. "Shelfware" is the New Churn Risk

Distribution does not equal adoption, and adoption does not guarantee pricing power. The market is treating generic AI features as a commodity. This is seen in two ways: the "shelfware risk" facing tools like Microsoft 365 Copilot, where low (<2%) adoption fails to justify a premium add-on, and the defensive bundling of companies like Zoom, who are forced to give AI away for free to reduce churn. For product leaders, this is a critical warning: features not deeply embedded into essential workflows will fail to create the value required to drive new revenue.

Key takeaway:

The competitive edge in 2026 will not come from having the best LLM, but from having the most resilient business model. The market is shifting away from generic tools for humans and toward specialized infrastructure for agents. Operators must now re-align their product and pricing strategies to sell outcomes, not just access.

The following table provides a top-level overview of the competing AI monetization strategies currently defining the enterprise software landscape.

Enterprise AI Monetization		
Company	AI Product	Monetization Model
--- Premium Add-Ons ---		
Microsoft	M365 Copilot	Premium Add-On (\$30/user/month; SMBs now at \$21/user/month)
Salesforce	Agentforce	Hybrid Add-On (Flex Credits + Per-user SKUs)
Google	Gemini for Workspace	Premium Add-On (Tiered pricing per user; \$20/user/month for Business, \$30/user/month for Enterprise)

Enterprise AI Monetization (continued)

Company	AI Product	Monetization Model
--- Tiered / Bundled Upgrades ---		
ServiceNow	Now Assist	Premium Tier Uplift (~60% premium for Pro Plus; consumption-based "assists" with starter packs)
SAP	Joule / Business AI	Cloud Bundle / Migration Driver (Included in RISE/Cloud tiers; Premium AI via tiered AI Units/user/month)
Workday	Illuminate	Bundled Value-Driver (Embedded + specific AI SKUs; Flex Credits for consumption scaling)
Atlassian	Rovo / Atlassian Intelligence	Forced Tier-Upgrade (Bundled into Premium/Enterprise plans; \$20/user/month annual)
--- Consumption / Outcome-Based ---		
Oracle	OCI Gen2 & Fusion AI	Vertical Integration (IaaS consumption + Bundled SaaS; per-character/token for models)
Snowflake	Cortex AI	Consumption Utility (Pay-per-credit; token-based for LLMs, GB-month for search)
MongoDB	Atlas Vector Search	Consumption (Compute/Storage on Atlas platform; included in cluster fees, no separate vector charge)
Asana	AI Teammates / AI Studio	Consumption (AI Credits with usage tiers; Basic included, Plus/Pro add-ons for advanced models)
UiPath	Autopilot	Hybrid Consumption (Per-bot / per-process pricing; unified platform units across services)

Enterprise AI Monetization (continued)

Company	AI Product	Monetization Model
--- Defensive / Commoditized ---		
Zoom	AI Companion	Defensive Bundle (Included for free with paid plans)
Adobe	Firefly	Hybrid Consumption (Credit Packs + Premium Tiers; \$9.99/month Standard for 2,000 credits)

Table of Contents

Executive Summary	2
The Outcome-Based Pivot	7
The Infrastructure Moat	10
The Incumbent's Trap	13
The Commoditized Middle	15
Implications for Your Role	17
Conclusion	19
About ARPU	21

The Outcome-Based Pivot

Architecting the Post-Seat Business Model

The most significant divergence in the 2025 software market is not technological, but commercial. As the industry confronts the limitations of the per-user license in an age of automation, two distinct monetization frameworks have emerged as successful alternatives.

As the market's definitive Monetizers, **ServiceNow** and **Snowflake** have effectively decoupled their AI revenue growth from simple headcount expansion. They have achieved this by deploying pricing architectures that align vendor revenue directly with the value of the automated work performed. These two models—the Premium Tier Uplift and the Consumption Utility—represent the most prominent and successful frameworks for monetizing enterprise AI today.

ServiceNow's Premium Tier Uplift

ServiceNow's success in monetizing AI comes from a deep understanding of its own strategic advantage: for thousands of enterprises, ServiceNow is the system of record for mission-critical workflows like IT and customer support. Instead of trying to sell a new, separate AI tool, they simply offered a more powerful version of the platform their customers already depended on.

The Mechanism: An Upgrade, Not an Add-On

The core of this strategy is the "Pro Plus" SKU. This is best understood not as an "AI feature," but as a premium upgrade package.

Imagine buying a car. The standard model gets you from A to B. The premium package costs more but includes advanced driver-assist features that handle traffic and parking for you. ServiceNow applied this logic to their software. Customers were already on the "Pro" tier. By upgrading to "Pro Plus" for a premium—estimated by analysts at approximately 60%—they unlock a suite of embedded AI agents. These agents aren't in a separate chat window; they work in the background, autonomously handling tasks like summarizing case notes or resolving common IT tickets.

Strategic Advantages

This model succeeds because it reframes the purchasing decision from a speculative bet on "AI" to a simple, quantifiable investment in operational efficiency.

- **For the Customer:** The value proposition is "upgrade your platform and achieve 30% greater workflow efficiency." An IT leader can easily justify paying a 60% license premium if it means they can handle a 30% increase in ticket volume without hiring more staff.
- **For the Vendor:** This model removes procurement friction. The sales team isn't selling a new, unproven product that requires a separate budget and risk review. They are selling an upgrade to an existing, mission-critical system. This protects the core revenue stream while capturing the value of AI automation through a higher average selling price (ASP).

Market Validation

The financial data confirms the efficacy of this approach. In the third quarter of 2025, ServiceNow reported closing 103 deals with net new Annual Contract Value (ACV) exceeding \$1 million. Furthermore, the company's generative AI products reached a revenue run-rate surpassing \$500 million, making it the fastest-growing new product line in the company's history. This performance validates that enterprises are willing to pay a significant premium for AI when it is packaged as a workflow upgrade rather than a standalone tool.

Snowflake's Consumption Utility

Snowflake has executed a strategy that is fundamentally different from its SaaS peers, treating AI not as a software license but as a metered utility that runs on the enterprise's most critical asset: its data. This model leverages the company's existing credit-based architecture to monetize AI adoption without the friction of a per-seat license.

The Mechanism: Metered Intelligence

Under this model, AI capabilities such as Cortex are accessed via the consumption of Snowflake credits. There is no per-user subscription fee. Customers pay only for the compute resources their AI workloads utilize, scaling their spend linearly with usage. This includes token-based pricing for LLMs and storage-based pricing for vector search, aligning costs directly with consumption.

Strategic Advantages

Snowflake's strategy is built on a powerful paradox: it uses low technical lock-in to create immense operational data gravity:

- **Inverted Sales Motion:** By eliminating the high fixed cost of a per-seat add-on, Snowflake transforms AI from a high-friction "upsell" into a low-friction "acquisition" tool. This lowers the barrier to entry, encouraging widespread experimentation.
- **Data Gravity:** Snowflake's platform was designed to be cloud-agnostic and support open formats, making it easy for enterprises to consolidate their data in one place. Once a company's data is centralized, it develops "gravity"—attracting more tools, pipelines, and workflows. For AI, this data gravity is a huge advantage. As enterprises build AI applications, their entire development process naturally orbits this data foundation. By owning the data layer, Snowflake inherently owns the AI workflow without needing to own the end-user's business process.

Market Validation

The financial data confirms that this data-centric, consumption-based model is a powerful growth engine.

Snowflake reports that AI capabilities now power 25% of all deployed use cases and influence approximately 50% of new customer wins, with over 6,100 accounts using its AI features weekly. This usage-based revenue stream contributed to a product revenue acceleration to 32% year-over-year in Q2 FY2026.

This performance validates that a consumption model is highly effective when the vendor owns the foundational layer—in this case, the data—where value is naturally consumed and scaled.

Implication for Operators

The 2026 budgeting cycle has become an existential inflection point for product and GTM leaders. Roadmaps that remain anchored to per-seat chatbot add-ons (like a Copilot) or generic assistants are now structurally misaligned with proven customer willingness-to-pay. The new industry benchmark, set by ServiceNow and Snowflake, is clear: AI value must be captured either through a forced platform-tier uplift that re-anchors the customer relationship or through immediate transition to metered consumption.

The Infrastructure Moat

Vertical Integration as the Ultimate Differentiator

While software vendors experiment with pricing models to capture value, a parallel competitive dynamic has emerged in the infrastructure layer. The scarcity of high-performance compute—specifically GPUs and specialized data centers—has created a unique opportunity for vendors who can bridge the gap between software applications and physical hardware.

Oracle's Infrastructure Play

Oracle serves as the definitive case study for this strategy. By successfully pivoting from a legacy database vendor to a vertically integrated cloud infrastructure provider, Oracle has built a competitive moat that pure-play SaaS vendors cannot replicate.

Owning the Compute

In a typical SaaS model, the vendor builds the application and rents the underlying infrastructure from a hyperscaler (AWS, Azure, Google Cloud). This dependency caps margins and leaves the application vendor vulnerable to infrastructure constraints.

Oracle inverted this dynamic. By building Oracle Cloud Infrastructure (OCI), a massive, capital-intensive network of Gen2 data centers, the company positioned itself to capture the primary capital expenditure boom of the AI era.

Strategic Advantages

Oracle's model has a few distinct advantages:

- **Capturing the "Training" Spend:** While most software vendors fight for "inference" revenue (running the model), Oracle's infrastructure captures the massive upfront "training" revenue. The company successfully positioned OCI as the premier "spillover" utility for a capacity-constrained market. This strategy was validated when Microsoft itself partnered with Oracle to offload massive OpenAI training workloads to OCI, proving that Oracle's network topology can compete directly with the hyperscalers.

- **Private Data + Public-Scale Models Without Trade-offs:** Pure SaaS players force a painful choice: move sensitive data to public models or keep it on-premise and lose access to the latest LLMs. Oracle eliminates that dilemma entirely. Because its applications and database sit natively inside OCI, every enterprise record is instantly addressable by the latest hosted LLMs without ever leaving a private, auditable boundary. This turns data gravity from a liability into a direct revenue accelerant—enterprises can finally run frontier models on their crown-jewel data without rewriting pipelines or accepting governance risk.
- **Sovereign AI:** Oracle differentiated its offering by building "Sovereign Cloud" regions—physically and logically isolated data centers designed to meet the strict data residency requirements of governments and regulated industries. This allows them to win high-value public sector contracts that generalist clouds often miss.

Market Validation

The financial results of 2025 validate the success of this vertical integration strategy. While growth in Oracle's traditional application business remains steady, its infrastructure business has exploded.

- **Revenue Acceleration:** OCI revenue growth accelerated to over 52% yoy, significantly outpacing the growth rates of many pure-play software competitors.
- **Demand Signals:** The intensity of demand is evidenced by over 300% spike in AI GPU consumption. This indicates that the market is treating Oracle not just as a software vendor, but as a critical utility provider for the AI economy.

Implication for Operators

For pure-play application vendors that continue to rent compute from third-party hyperscalers, the strategic window has effectively collapsed from years to months. Only two defensible long-term positions remain: (1) secure ownership or deep control of the infrastructure/data layer through partnership or acquisition, or (2) achieve such extreme workflow entrenchment that the product becomes acquisition-proof for the infrastructure giants. The middle ground—remaining a horizontal application layer on top of AWS, Azure, or GCP—has become the single most exposed position in the post-2025 stack.

“

Do you think JPMorgan Chase makes all of their internal data publicly available? Companies want to be able to use AI models on top of their own data. That is essential. Oracle applications make all of the [internal] data available to AI models... We have all of these LLMs in the Oracle cloud. This is our value proposition.

- Larry Ellison, Oracle Founder & CTO

Q4 2025 Earnings Call

The Incumbent's Trap

Navigating the Friction of Scale and Legacy Models

While agile players execute new pricing and infrastructure strategies, the industry's software giants—**Salesforce** and **Microsoft**—face a more complex challenge. Their massive installed bases, which typically serve as a competitive moat, have paradoxically created inertia. These giants are currently attempting to pivot from feature Bundlers to outcome Monetizers before their core pricing power erodes

But there are two "traps" in this pivot: the difficulty of transitioning pricing models at scale, and the risk of "shelfware" in general-purpose deployments.

Salesforce's Pricing Pivot

Salesforce is actively managing a high-stakes transition from its traditional "per-user" model to a new consumption-based framework for its Agentforce platform. This pivot highlights the extreme difficulty of changing the commercial terms of a decades-old relationship with the enterprise buyer:

- **The Friction Point:** Initial attempts to price agents on a per-conversation basis encountered significant resistance. The original \$2 per conversation model drew backlash for its unpredictability and high costs, particularly among mid-market organizations accustomed to predictable subscription fees, leading to heavy adoption hurdles.
- **The Strategic Correction:** In response, Salesforce introduced "Flex Credits" in May 2025 (priced at \$0.10 per action, with packs like 100,000 credits for \$500 in some bundles). This hybrid mechanism—now including unmetered employee-facing usage in new Agentforce 1 Editions and add-ons launched in Summer 2025—attempts to bridge the gap, giving customers the predictability of pre-purchased credits while allowing Salesforce to monetize variable agent work. Flex Agreements further enable swapping seats for credits, adapting to evolving AI strategies.
- **Operational Reality:** This transition requires a fundamental retraining of the sales force. Reps who have spent their careers selling seats must now learn to sell consumption capacity and outcomes, a shift that often creates short-term execution drag. Despite these challenges, Agentforce adoption is gaining traction, with over 12,500 deals closed since launch and more than 6,000 converted to paid arrangements as of Q2 FY2026, contributing to AI and Data Cloud ARR surpassing \$1.2 billion.

Microsoft's Shelfware Risk

Microsoft's challenge is not pricing, but adoption depth (here we are referring to M365, not its Azure business). With its "Copilot" strategy, Microsoft achieved massive distribution overnight, embedding AI into the world's most popular productivity suite. However, distribution has not immediately translated into habitual usage:

- **The Data Reality:** Leaked internal metrics suggest that adoption of Microsoft 365 Copilot among eligible commercial seats remains low, estimated at below 2%.
- **The "Generalist" Problem:** The low adoption rate signals a "shelfware" risk. General-purpose assistants that sit alongside the workflow (e.g., a chatbot in Word) often fail to become essential because they require the user to actively change their behavior.
- **Operational Reality:** This creates a downstream retention risk. CIOs who purchased pilot licenses may churn if they do not see widespread employee adoption before the renewal cycle. It underscores the reality that for AI to stick, it must be embedded into specific, high-value workflows, not just available as a general tool.

Implication for Operators

The struggles of these incumbents provide a clear lesson: scale is not a strategy. Simply having a large customer base or a powerful brand does not guarantee AI monetization. Success requires aligning the business model with the customer's risk tolerance (pricing) and ensuring the product solves a specific, unavoidable problem (workflow integration). As Agentforce and Copilot evolve toward agentic capabilities, operators must prioritize measurable outcomes over broad enablement to avoid the traps of inertia and underutilization.

The Commoditized Middle

The High Cost of Low Stickiness: When Usage Doesn't Equal Value

The final competitive shift defining the 2025 market is the evolution of Shelfware Risk. Historically, shelfware meant software that wasn't used. In the AI era, it means software that is used passively but fails to deepen the customer's dependence on the platform.

Zoom and **Adobe** illustrate this danger. Both have achieved high AI adoption, yet both face valuation compression because their AI features have not created the "workflow stickiness" required to drive pricing power.

Zoom's "Defensive" Trap

Zoom serves as the archetype for The Bundler. It has executed a strategy of aggressive bundling, including its "AI Companion" at no additional cost. While successful as a defensive moat, it has trapped the company in a low-growth cycle:

- **The Strategy:** By bundling AI, Zoom effectively bought retention. The strategy worked: churn hit a record low of 2.7% in Q4 Fiscal 2025.
- **The Shelfware Reality:** However, the AI usage is largely passive (e.g., meeting summaries) rather than agentic (executing workflows). Because the AI acts as a post-meeting aid rather than a real-time operator, it creates value but not dependence.
- **The Consequence:** AI has become a permanent cost center (pressuring gross margins to ~76%) that protects the core but adds minimal net-new growth. Revenue growth remains in the low single digits (~3-4%), proving that without a premium SKU, AI is merely an infrastructure tax.
- **The Operator's Lesson:** Free is a trap. If your AI strategy is simply "doing the same thing, but faster," you cannot charge for it. To escape the defensive trap, AI must enable *new* workflows—like the forthcoming \$12 custom AI add-on—that were previously impossible.

Adobe's Bypass Risk

Adobe faces a different flavor of shelfware risk: the Bypass Risk. Despite solid financials (11% revenue growth), the company's stock has compressed ~26% YTD due to market fears that Generative AI allows non-experts to bypass complex professional tools entirely:

- **The Stickiness Problem:** Adobe's Firefly model is a leader in commercial safety, generating over 29 billion assets in Q3 FY2025. However, the workflow often remains exportable—users generate an image in Firefly and then leave the ecosystem to assemble it in Canva or PowerPoint.
- **The Consequence:** This creates an ideation-versus-execution gap. While Adobe dominates the high-end professional workflow, it risks becoming shelfware for the 40% of the market comprised of casual users who find "good enough" results elsewhere. AI-influenced ARR now exceeds \$250 million (ahead of FY2025 target), but analysts caution material Firefly revenue impact may not hit until 2026–2027.
- **The Operator's Lesson:** High usage does not equal high stickiness. If your AI output can be easily exported to a cheaper competitor, you lack a defensible moat. The goal is to embed the AI so deeply into the end-to-end *process* that leaving the platform becomes operationally painful.

Implication for Operators

The experiences of Zoom and Adobe highlight a critical danger: AI features are rapidly becoming table stakes. To avoid the commoditized middle, operators must ensure their AI strategy is built on deep workflow integration—where the AI executes work, not just generates content—to create the lock-in required for long-term monetization. As bundling gives way to agentic enablement, the winners will be those who turn passive tools into indispensable infrastructure.

Implications For Your Role

Major market events create cascading effects that ripple across the entire ecosystem, impacting every business and function in a unique way.

The shift away from seat-based pricing is not just a financial trend; it's the first domino. The cascading effects include a complete re-evaluation of product roadmaps, a fundamental restructuring of go-to-market teams, and a new framework for strategic M&A.

While most analysis focuses on the initial event, ARPU specializes in mapping these cascading effects directly to your specific role and business model.

To demonstrate, let's deconstruct the single most important strategic challenge from this report—The Inevitable Collapse of the Per-Seat Business Model—and show how it creates distinct, actionable intelligence for different leaders across a B2B software company.

One Trend, Three Bespoke Implications: Monetizing After the "Seat"

The slow death of the per-user license demands a new operational playbook. The value you create is no longer tied to the number of humans using your software, but to the amount of work your software performs.

For Head of Product at application-layer SaaS vendors:

Your roadmap is now a monetization engine. "Productivity features" are no longer enough; you must build *meterable* events that can be tied to a consumption or outcome-based SKU. The strategic priority must shift from building assistive "Copilots" (which are hard to price) to building autonomous "Agents" that execute core workflows (which are easy to price). Workflow stickiness is now your most important product metric.

The Key Question: Does your 2026 roadmap prioritize building specific, transactional agents that can justify a consumption-based revenue stream, or are you still building general-purpose assistants that will be commoditized?

For CRO / VP of Sales at seat-based B2B SaaS:

Most enterprise sales organizations remain optimized for a business model that is already dying. The classic “land-and-expand” playbook—built on large upfront bookings, multi-year seat commitments, and quotas tied to initial ACV—systematically punishes the exact behaviors required to win in a consumption/outcome world. Successful consumption motions demand post-sale cultivation, technical co-architecture, and expansion credit for usage growth, not Day-1 bookings.

The Key Question: Are sales compensation plans still incentivizing the hunters of a seat-based era while the market has permanently shifted to rewarding cultivators of usage and outcomes?

For Head of Strategy / Corp Dev at public SaaS platforms:

The traditional framework for valuing SaaS is at risk of becoming obsolete. A target's seat count or ARR growth from seat expansion are no longer reliable indicators of future success; they are vanity metrics masking a brittle business model. The most valuable acquisition targets are no longer those with the biggest user bases, but those who have successfully cracked the code on vertical AI and outcome-based pricing in a specific niche.

The Key Question: Are you overvaluing your own business and your acquisition targets by using outdated “per-seat” TAM calculations that ignore the deflationary pressure of AI agents?

Ready to see what these trends mean for your 2026 roadmap?

Request a Complimentary, Personalized Sample Brief

Contact us online at: <https://arpu.hedder.com/contact-us/>

Or email our team at: arpu@hedder.com

Conclusion

The Next Phase of AI Monetization

2025 settled the question of who can actually charge for AI and who cannot.

The results are now visible in the numbers:

- Companies that own mission-critical workflows (ServiceNow) or the underlying data/compute (Snowflake, Oracle) are successfully adding new AI revenue streams without cannibalizing their core.
- Everyone else is either bundling AI defensively, pivoting pricing under pressure, or watching adoption stall.

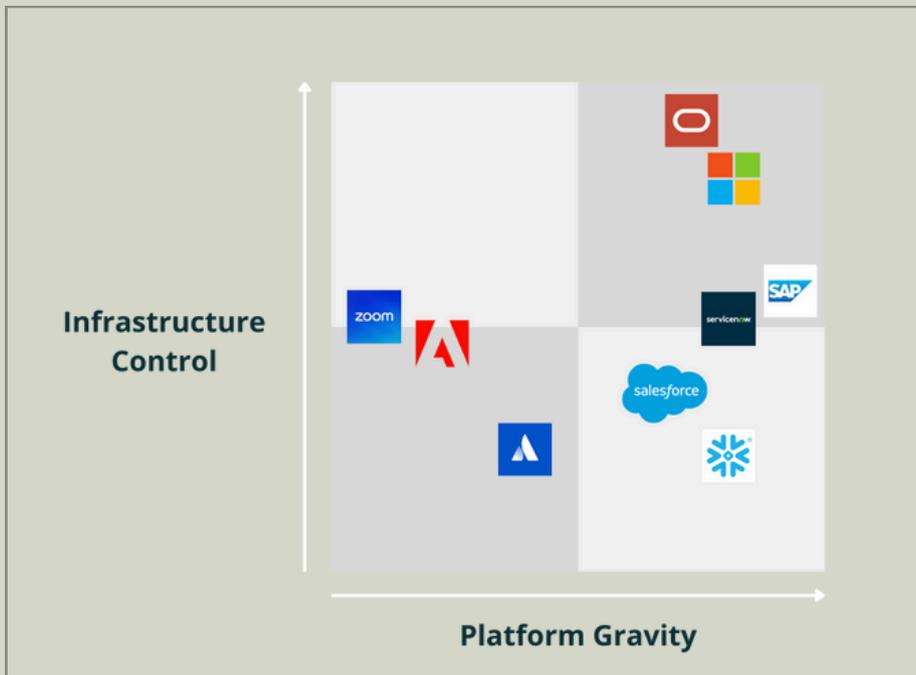
The clearest trend heading into 2026 is the accelerating move away from per-seat pricing for any AI-heavy product. Successful AI deployments reduce headcount or automate seats; the more successful the AI, the faster the old pricing model breaks. We expect most incumbent SaaS vendors to introduce some form of consumption, outcome, or tier-uplift pricing in the next 12–18 months simply to stay revenue-neutral.

The durable advantage no longer belongs to the best model or the prettiest interface. It belongs to two places only:

1. The layer that already owns the enterprise data (Oracle, Snowflake)
2. The layer that already executes the enterprise workflow (ServiceNow, a few vertical leaders)

Everything sitting in between—horizontal assistants, add-on copilots, exportable generative tools—will face steady price compression unless it moves decisively into one of those two layers.

The entire 2025–2026 landscape can be reduced to a single framework (next page):



The matrix reveals the two defensible paths to winning in the AI era: high Infrastructure Control (the top half) or high Platform Gravity (the right half). Platform Gravity measures a system's entrenchment as the indispensable center of a mission-critical business process (like SAP/ServiceNow) or the enterprise data foundation (like Snowflake). The strongest competitive position belongs to the vertical giants in the top-right who possess both. The platform players in the bottom-right have an equally powerful moat built on deep operational entrenchment. Companies in the bottom-left Danger Zone lack a clear source of gravity and face significant commoditization risk.

The practical question for every leadership team in 2026 planning season is therefore very simple:

Do we own the data or the workflow well enough to charge for the work our AI performs—or are we still trying to charge by the seat for something that is actively removing seats?

That single question will separate next year's winners from everyone else.

About ARPU

ARPU provides bespoke, function-specific intelligence for tech operators. In a market saturated with generic research and noise, we deliver signal.

Our analysis translates market-wide events and trends into intelligence briefs tailored to the specific challenges and questions of your role. We exist to help you make better, faster decisions.

Go Deeper.

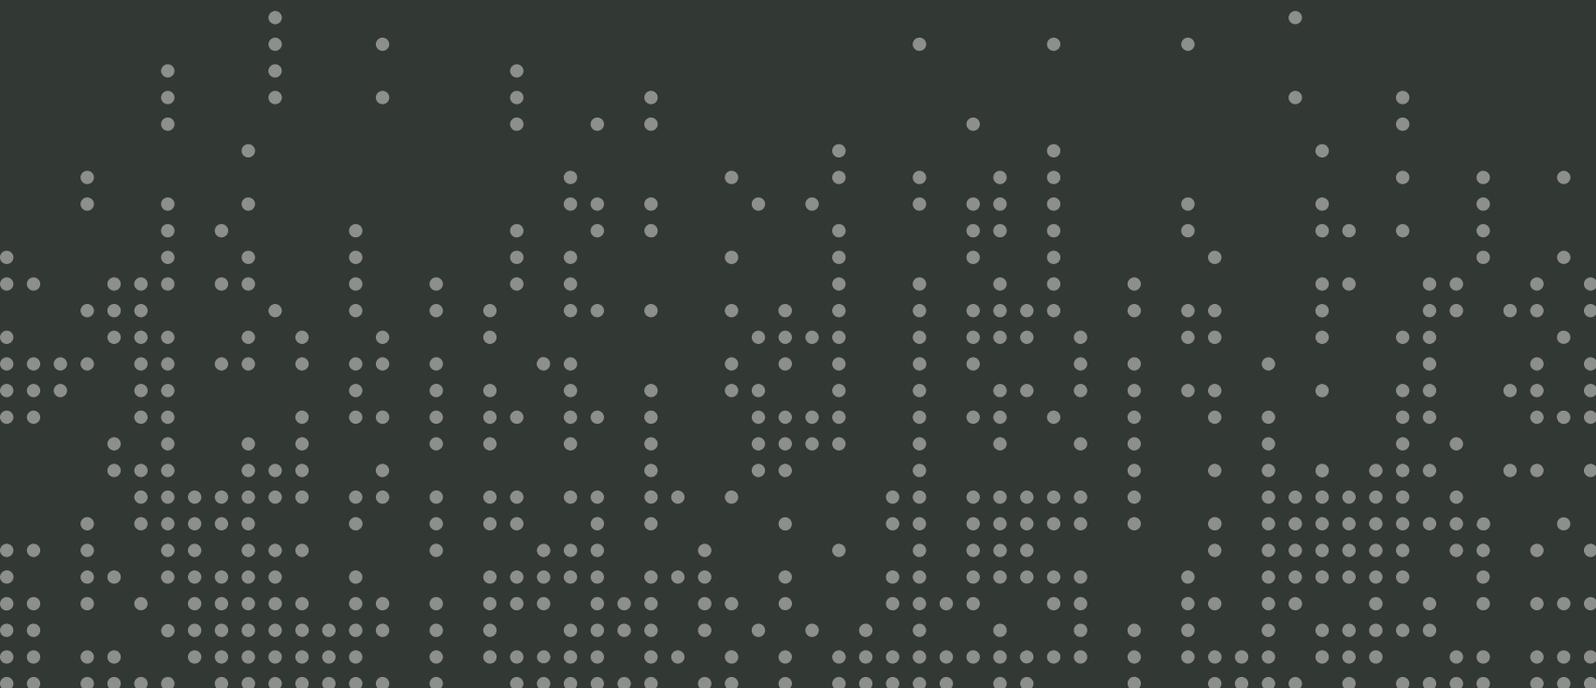
This report provided a market-wide view of the enterprise AI monetization landscape. But your organization, your products, and your strategic questions are unique.

Ready to see what these trends mean for your 2026 roadmap?

Request a Complimentary, Personalized Sample Brief

Contact us directly at arpu@hedder.com

Or visit our website at <https://arpu.hedder.com/>



Sources

Our analysis is based on a review of the following sources, in addition to the quarterly earnings call transcripts and presentations for the companies covered in this report.

Agentforce Pricing Update: Salesforce Announces Major Changes
(<https://www.salesforceben.com/agentforce-pricing-update-salesforce-announces-major-changes/>)

AI Powering 25% of Snowflake Customer Projects; Revenue Jumps 32% - Cloud Wars
(<https://cloudwars.com/ai/ai-powering-25-of-snowflake-customer-projects-revenue-jumps-32/>)

Distributed Cloud Services | Oracle (<https://www.oracle.com/cloud/distributed-cloud/>)

Enable Secure and Scalable Self-Service Platforms for Generative AI and LLMs within OCI
(<https://docs.oracle.com/en/solutions/oci-generative-ai-llm-platforms/index.html>)

Flexible Pricing is a Full Circle Moment for Salesforce (<https://www.salesforce.com/blog/flex-credits/>)

From \$2 Conversations to \$0.10 Actions: Inside New Agentforce Pricing (<https://www.concret.io/blog/new-agentforce-pricing-model>)

Inside Salesforce's Retraining Strategy Amid Fears of AI Job Loss (<https://www.newsweek.com/nw-ai/inside-salesforces-retraining-strategy-amid-fears-of-ai-job-loss-10993019>)

Introducing Oracle AI Database 26ai: Next-Gen AI-Native Database for All Your Data
(<https://blogs.oracle.com/database/oracle-announces-oracle-ai-database-26ai>)

Key Insights From the 2025 Microsoft 365 and Copilot Survey
(<https://www.gartner.com/en/documents/6548002>)

Microsoft 365 Copilot adoption report (<https://learn.microsoft.com/en-us/viva/insights/advanced/analyst/templates/microsoft-365-copilot-adoption>)

Microsoft 365 Copilot impact report (<https://learn.microsoft.com/en-us/viva/insights/advanced/analyst/templates/microsoft-365-copilot-impact>)

Microsoft 365 Copilot's commercial failure (<https://www.perspectives.plus/p/microsoft-365-copilot-commercial-failure>)

OpenAI to Use OCI for AI Training, Oracle to Build "Very Large Data Center"
(<https://www.datacenterdynamics.com/en/news/openai-to-use-oci-for-ai-training-oracle-to-build-very-large-data-center/>)

OpenAI, Oracle, and SoftBank Expand Stargate with Five New AI Data Center Sites
(<https://openai.com/index/five-new-stargate-sites/>)

Oracle AI Database 26ai Powers the AI for Data Revolution | Morningstar
(<https://www.morningstar.com/news/pr-newswire/20251014sf97015/oracle-ai-database-26ai-powers-the-ai-for-data-revolution>)

Oracle and NVIDIA to Deliver Sovereign AI Worldwide (<https://nvidianews.nvidia.com/news/oracle-nvidia-sovereign-ai>)

Oracle Embeds Generative AI Across the Technology Stack to Enable Enterprise AI Adoption at Scale (<https://www.prnewswire.com/news-releases/oracle-embeds-generative-ai-across-the-technology-stack-to-enable-enterprise-ai-adoption-at-scale-302041444.html>)

Powering Smarter Insights, Scalable AI on OCI (<https://www.oracle.com/ai-data-platform/ai-data-platform-ai-transformation-oci/>)

Salesforce Shifts to Agentforce Flex Credits: Addressing Adoption Barriers, But Challenges Remain (<https://salesforcedevops.net/index.php/2025/05/15/salesforce-shifts-to-agentforce-flex-credits/>)

ServiceNow CFO outlines method to Pro Plus SKU pricing | Constellation Research Inc. (<https://www.constellationr.com/blog-news/insights/servicenow-cfo-outlines-method-pro-plus-sku-pricing>)

ServiceNow Pricing Guide: Custom Quotes for Tailored IT Solutions (<https://www.rezolve.ai/blog/servicenow-pricing-guide-custom-quotes-for-tailored-it-solutions>)

Snowflake Cortex AI Functions (including LLM functions) | Snowflake Documentation (<https://docs.snowflake.com/en/user-guide/snowflake-cortex/aisql>)

Snowflake Cortex versus Other LLMs: Pricing, Tradeoffs, Benchmarks, and When to Choose Each | Keebo (<https://keebo.ai/2025/10/17/snowflake-cortex-versus-other-llms/>)

Snowflake Q2 Revenue Jumps 32 Percent | Nasdaq (<https://www.nasdaq.com/articles/snowflake-q2-revenue-jumps-32-percent>)

Salesforce Agentforce pricing models continue to evolve (<https://martech.org/salesforce-introduces-new-agentforce-pricing-models/>)

Salesforce Agentforce Pricing Update (Q3 2025): Flex Credits, Agreements, and Editions (<https://aqualabs.com/blog/agentforce-pricing-gets-a-long-overdue-fix-flex-credits-are-now-live/>)

Salesforce Bets Big on Agentforce: Will Adoption Drive Revenues? (<https://www.nasdaq.com/articles/salesforce-bets-big-agentforce-will-adoption-drive-revenues>)

Salesforce Introduces New Flexible Agentforce Pricing to Accelerate the Digital Labor Revolution (<https://www.salesforce.com/news/press-releases/2025/05/15/agentforce-flexible-pricing-news/>)

Salesforce issues weak revenue guidance even as earnings beat estimates (<https://www.cnbc.com/2025/09/03/salesforce-crm-q2-earnings-report-2026.html>)

Salesforce Makes Changes to Its Agentforce Pricing Model (Again!) (<https://www.cxtoday.com/crm/salesforce-makes-changes-to-its-agentforce-pricing-model-again/>)

Salesforce revamps Agentforce pricing with Flex Credits: What you need to know (<https://www.constellationr.com/blog-news/insights/salesforce-revamps-agentforce-pricing-flex-credits-what-you-need-know>)

Snowflake's AI-Driven Growth and Earnings Momentum: Why It's a Must-Watch AI Stock in 2025
(<https://www.ainvest.com/news/snowflake-ai-driven-growth-earnings-momentum-watch-ai-stock-2025-2508/>)

Understanding cost for Cortex Search Services | Snowflake Documentation
(<https://docs.snowflake.com/en/user-guide/snowflake-cortex/cortex-search/cortex-search-costs>)

What's new in Microsoft 365 Copilot | July 2025
(<https://techcommunity.microsoft.com/blog/microsoft365copilotblog/what%25E2%2580%2599s-new-in-microsoft-365-copilot--july-2025/4438253>)

What Is Sovereign AI? (<https://www.oracle.com/asean/artificial-intelligence/what-is-sovereign-ai/>)

Why Oracle Stands Out for LLM Deployment Over AWS, SAP, and Google Cloud
(https://medium.com/@sales_96306/why-oracle-stands-out-for-llm-deployment-over-aws-sap-and-google-cloud-eecc48093cfb)