RESEARCH BRIEF

The Unmanaged Resolution Economy:

Sizing AI Agent Task Completions Across Enterprise and SMB Markets, 2026–2030

> ““
>
> *No sheriff in sight. Billions of agentic completions are occurring daily with no standard definition, no governance framework, and no pricing discipline.*
>
> — The Unmanaged Resolution Economy, CPAG Research Brief, May 2026

Data current as of May 2026

Sources: Gartner, Salesforce, Intercom, Zendesk, Epoch AI, McKinsey, IDC, Forrester, SBA

A NOTE ON METHODOLOGY AND CONFIDENCE

No industry-standard definition of an AI agent task completion currently exists. This report constructs estimates by triangulating across three data tiers: (1) vendor-disclosed, explicitly billed resolution counts; (2) verifiable AI task completions reported in earnings disclosures and analyst surveys but not billed as resolutions; and (3) estimated untracked agentic completions derived from adoption rates, workforce data, and task frequency proxies. Each tier carries a different confidence level, stated explicitly throughout. Aggregate figures represent directional ranges, not audited totals. Enterprise is defined as organizations with 500 or more employees. SMB is defined per the U.S. Small Business Administration standard of fewer than 500 employees.

Executive Summary

AI agents are completing billions of discrete tasks inside enterprise and SMB operations every day in 2026. Customer service inquiries are being resolved. Contracts are being summarized. Code is being reviewed. Finance exceptions are being flagged and cleared. Support tickets are closing without a human agent touching them. The volume of agentic task completions is growing at a rate that is structurally reshaping enterprise software economics.

None of it is being measured in a consistent way.

Vendors are inventing proprietary units. Salesforce introduced the Agentic Work Unit (AWU) in its Q4 FY26 earnings disclosure, reporting 2.4 billion AWUs delivered to date. Intercom counts resolutions at $0.99 each. Zendesk bills automated resolutions at $1.50 to $2.00 per event. Every major platform has a different definition of what constitutes a completed unit of agentic work. No two are compatible. None fully satisfy the three criteria that make a unit commercially durable: verifiable, attributable, and finite.

This report sizes the agentic completion economy as it actually exists today, not as vendors report it, and forecasts its trajectory through 2030. The market is large, growing rapidly, and operating without a definitional or governance infrastructure adequate to support its scale. That gap is not a philosophical problem. It is a commercial liability accumulating on every enterprise balance sheet that has deployed AI agents without defining what they are doing.

Key findings:

- Estimated 86 billion AI agent task completions occurred across US enterprise and SMB markets in 2025. By 2030, that figure reaches between 1.1 trillion (plateau scenario) and 4.2 trillion (accelerated scenario) annually.
- Only a small fraction of these completions are priced as resolutions. CPAG estimates that fewer than 3% of enterprise agentic task completions in 2026 carry any outcome-based pricing. The remaining 97% are absorbed into flat-rate seat contracts or unmonitored infrastructure costs.
- The largest enterprise software vendor on the planet invented its own measurement unit rather than adopt a standard, because no standard exists. Salesforce's Agentic Work Unit (AWU) is commercially unverifiable, definitionally broad, and platform-proprietary.

- Inference costs are falling dramatically but enterprise AI bills are rising. Gartner projects a 90% reduction in LLM inference costs by 2030. Yet enterprise AI spend is rising because agentic workflows consume 10–50x more tokens per task than single-query interactions.
- 21% of organizations have a mature governance model for autonomous AI agents. Gartner projects that more than 40% of agentic AI projects will be cancelled by 2027. The primary cause is not model quality. It is the absence of audit trails, attribution frameworks, and outcome definitions.
- SMB adoption is closing the enterprise gap faster than any prior technology cycle. The SBA reports the large-to-small adoption ratio narrowed from 1.8x to 1.2x between 2024 and 2025.

Section 1: Defining the Resolution for Measurement Purposes

Before any market can be sized, the unit of measurement must be defined. In AI agent markets as they exist today, that definition does not exist in any standard form. This is not a minor methodological inconvenience. It is the central structural problem the market must solve before agentic AI can be priced, audited, or governed at scale.

The industry has produced a proliferation of competing terms. Salesforce measures Agentic Work Units. Intercom counts resolutions. Zendesk tracks automated resolutions and deflection rates as separate metrics, noting explicitly that deflection and resolution are not the same event. Microsoft Copilot reports assisted actions. ServiceNow measures flow executions. None of these definitions are compatible.

> ““
>
> *“Salesforce introduced the Agentic Work Unit (AWU), defining it as ‘discrete tasks executed by AI agents in production across the Salesforce platform, including Slack.’ 2.4 billion AWUs were reported as of Q4 FY26. The unit is the most prominent attempt in enterprise software to measure agentic completions at scale. It is also commercially unverifiable. An AWU cannot be priced, disputed, or audited by a CFO because it encompasses tasks of radically different complexity, cost, and business value within the same definition.”*

AWU illustrates the problem precisely. It is a disclosure metric, not a commercial unit. The fact that the largest enterprise CRM company on the planet invented a proprietary unit rather than adopting or proposing a standard is the clearest evidence available that no such standard exists.1

The Three-Tier Taxonomy Used in This Report

Because no single industry definition exists, this report constructs estimates using a three-tier taxonomy designed to reflect actual market conditions while being transparent about confidence levels at each tier.

Tier	Definition	Examples	Confidence	Est. Share of Volume
Tier 1 Vendor disclosed, explicitly billed resolutions	Completions for which a specific per-resolution or per-unit fee is charged and disclosed by the vendor	Intercom Fin ($0.99/resolution), Zendesk AI ($1.50–$2.00/automated resolution)	HIGH	~3%
Tier 2 Verifiable completions, not billed as resolutions	AI agent task completions measurable via vendor metrics but absorbed into seat contracts or platform fees rather than priced per event	Salesforce AWUs, Microsoft Copilot assisted actions, ServiceNow flow executions, AI deflection rates in CX surveys	MODERATE	~22%
Tier 3 Estimated untracked agentic completions	AI-driven task completions occurring inside enterprise and SMB operations with no attribution, no measurement, and no pricing. The dark matter of the resolution economy.	Background agent loops, unlogged automation runs, AI completions inside bundled copilot features, custom-built agents with no instrumentation	LOW	~75%

METHODOLOGICAL NOTE

The Tier 3 estimate is derived from the gap between total agentic AI deployment rates (Gartner: 40% of enterprise applications will embed task-specific agents by end of 2026, up from under 5% in 2025) and Tier 1 plus Tier 2 measured volumes. The 75% estimate is directional. The actual share is unknowable by definition, since untracked completions are untracked.

What a Durable Resolution Unit Requires

For a resolution to function as a commercial pricing unit rather than a marketing disclosure metric, it must satisfy three criteria. These criteria are not arbitrary. They are the minimum requirements for a unit that can appear in a contract, survive a billing dispute, and be audited by an enterprise procurement committee.

Verifiable: The completion must be objectively confirmable against a defined state. A ticket marked closed with no reopen within 48 hours is verifiable. A task that the AI “attempted” is not.

Attributable: The completion must be traceable to the AI agent's execution, not to a human override, a third-party process, or an ambiguous handoff. Attribution requires a complete action log.

Finite: The resolution must have a clear endpoint. Open-ended agentic loops that run indefinitely cannot be priced on a per-resolution basis without a defined scope boundary.

Zendesk's automated resolution metric comes closest among Tier 1 vendors to satisfying all three criteria, which is why it carries the highest per-unit price in the market ($1.50 to $2.00 versus Intercom's $0.99). Definitional rigor commands a price premium. The market is already demonstrating this.

Section 2: Market Sizing Methodology

The enterprise AI agent market does not publish aggregate resolution volume data. This report constructs estimates through triangulation across five input categories, each sourced from primary disclosures or tier-one analyst research. The methodology is stated in full because transparency is the only defense against the credibility attacks that aggregate estimates of this type attract.

Input Category	Primary Sources	Application in Model
Enterprise AI agent deployment rates	Gartner (August 2025), McKinsey Global AI Survey 2026, G2 Enterprise AI Agents Report	Baseline deployment penetration by company size and vertical
Disclosed vendor resolution volumes	Intercom (40M+ resolved conversations as of December 2025); Salesforce Q4 FY26 earnings (2.4B AWUs); Salesforce Q4 FY26 (3M support conversations in 12 months)	Anchor points for Tier 1 and Tier 2 volume estimates
Task completion rates and AI deflection benchmarks	Zendesk CX Trends 2025 (80%+ automated resolution claim); Gartner (agentic AI will autonomously resolve 80% of common customer service issues by 2029); First Page Sage Study (75.3% mean task completion rate, Q1 2026)	Resolution rate assumptions by use case vertical
Knowledge worker activity volumes	McKinsey Global AI Survey 2026 (6.4 median hours saved weekly per worker); US Bureau of Labor Statistics knowledge worker population estimates; SBA SMB workforce data	Task frequency and AI penetration by function for Tier 3 estimation
AI inference cost trajectories	Gartner (90% cost reduction by 2030); Epoch AI benchmark analysis (9x–900x price decline per year for quality-adjusted frontier models, 2024–2025); Statista (output cost per million tokens: ~$12 in 2022, under $2 by 2024)	Three-scenario cost model inputs

Section 3: Enterprise Market

Enterprise AI agent adoption crossed the production threshold in 2026. Gartner documents that 40% of enterprise applications will embed task-specific AI agents by end of 2026, up from under 5% in 2025. McKinsey reports that 23% of organizations are actively scaling agentic systems in at least one function. G2's August 2025 survey found 57% of companies had AI agents in production. The debate about whether enterprise AI agents are deployed is over. The debate about how to measure what they are doing has not yet begun.2

40%	2.4B	21%
Enterprise applications embedding task-specific AI agents by end of 2026, up from under 5% in 2025 *Gartner, August 2025*	Agentic Work Units delivered by Salesforce Agentforce to date as of Q4 FY26, growing 57% quarter-over-quarter *Salesforce Q4 FY26 Earnings, February 2026*	Organizations with a mature governance model for autonomous AI agents. 40% of agentic projects at risk of cancellation by 2027. *Gartner, 2026*

Volume Estimates by Use Case Vertical

Customer support is the most mature vertical for agentic task completions. Intercom's Fin agent has processed over 40 million resolved conversations as of December 2025, growing at a rate consistent with its $100M ARR trajectory and 350% year-over-year growth. Salesforce's own deployment of Agentforce on help.salesforce.com handled 3 million support conversations in 12 months, generating $100 million in annualized cost savings and reducing case volume by 8% year-over-year despite customer base growth. Reddit deployed Agentforce and deflected 46% of support cases, cutting average response time from 8.9 minutes to 1.4 minutes.3

Vertical	2026 Est. Annual Volume (US Enterprise)	Notes
Customer Support	~38 billion	Most mature vertical; anchored to Intercom and Salesforce disclosed volumes
IT Operations	~22 billion	Second largest by volume; significant Tier 3 untracked component
Sales & Marketing	~18 billion	Includes lead qualification, outreach, CRM updates
Finance & Accounting	~10 billion	Exception handling, reconciliation, reporting
HR & Recruiting	~7 billion	Screening, onboarding, documentation
Legal & Compliance	~4 billion	Contract review, compliance checks
Total (Enterprise)	~99 billion (range: 72B–135B)	CPAG mid-case estimate. Confidence: MODERATE.

Of those 99 billion estimated completions, CPAG estimates that fewer than 3 billion (approximately 3%) carry Tier 1 outcome-based pricing. The remaining 96 billion completions represent unpriced or underpriced agentic work being absorbed into legacy contract structures that were designed for human-paced workflows.

The Resolution Pricing Gap

The commercial implication is direct. If the average Tier 1 resolution price across the disclosed market is approximately $1.25 (between Intercom's $0.99 and Zendesk's $1.50 committed rate), then 96 billion unpriced enterprise completions represent a theoretical revenue gap of approximately $120 billion annually at current pricing levels. Not all of that is capturable under any realistic transition scenario. But even a 10% capture rate at a conservative $0.50 per resolution represents a $4.8 billion annual pricing opportunity that does not currently appear on any enterprise software vendor's revenue line.

> **INTERPRETIVE CAUTION**
>
> The $120 billion figure is a theoretical maximum, not a forecast. It assumes that every unpriced completion is equivalent in commercial value to a Tier 1 resolution, which is not true. Tier 3 completions in particular include background agent loops and low-complexity automations that would price at a fraction of a full resolution. The figure is presented to illustrate the order of magnitude of the pricing gap, not to project transition revenue.

Section 4: SMB Market

SMB AI agent adoption is accelerating at a rate that has closed the adoption gap with enterprise faster than any comparable technology cycle. The SBA Office of Advocacy reports the large-to-small business AI adoption ratio narrowed from 1.8x in February 2024 to 1.2x by August 2025. The US Chamber of Commerce documents small business AI investment rising from 36% in 2023 to 42% in 2024 to 57% in 2025. The primary drivers are the low cost of cloud-based AI tools and the disproportionate productivity impact of automation on small teams.4

57%	1.2x
SMBs investing in AI in 2025, up from 36% in 2023. Investment increased 58% in two years across businesses under 250 employees.	Current large-to-small business AI adoption ratio, down from 1.8x in February 2024. Gap is closing faster than any prior technology cycle.
business.com / Dialog SMB AI Survey, 2025	*SBA Office of Advocacy, August 2025*

Agentic AI specifically is entering the SMB market through the same customer support entry point that defined its enterprise trajectory, but with an 18 to 24 month lag. Intercom's standalone Fin product, available without the full Intercom platform at a 50-resolution monthly minimum, represents the primary Tier 1 channel for SMB resolution pricing today. Zendesk's per-resolution billing model, which began auto-charging overages in January 2026, is creating the first wave of SMB exposure to outcome-based pricing economics, with some accounts experiencing 3x to 4x their expected monthly costs as automation success exceeded budgeted resolution volumes.

The combined US enterprise and SMB total is estimated at approximately 117 billion AI agent task completions in 2026, with a wide confidence interval (83 billion to 163 billion) reflecting the structural data limitations in Tier 3 estimation. It is presented as a directional sizing of the unmanaged resolution economy, not an audited aggregate.5

Section 5: The Cost Trajectory, 2026–2030

The inference cost picture in 2026 contains a paradox that is creating significant confusion at the enterprise CFO level. Per-token prices are falling dramatically. Enterprise AI bills are rising. Both statements are true simultaneously, and the tension between them is where the governance argument gains its commercial urgency.

Gartner's landmark March 2026 forecast projects LLM inference costs will fall over 90% by 2030, making models up to 100 times more cost-efficient than comparable 2022 models. Epoch AI's analysis documents quality-adjusted token price declines of 9x to 900x annually for frontier models. Statista confirms the directional trend: average output costs per million tokens fell from approximately $12 for GPT-3.5 in 2022 to under $2 by 2024.6

Yet the same enterprises watching unit costs collapse are seeing monthly AI bills multiply. Three structural factors explain the divergence: agentic workflows are token-intensive, triggering 10 to 50 inference calls per task versus one for a simple chatbot query; background monitoring agents consume tokens continuously rather than on-demand; and vendors are absorbing inference cost savings into margin rather than passing them through to customers.

> ❝❝
>
> *"Even as Gartner projects inference costs will fall over 90% by 2030, the firm is explicit that these savings will not flow through to enterprise customers. The strategic implication: organizations without defined resolution units and per-resolution pricing discipline have no mechanism to capture the cost reduction that model efficiency is generating."*

Three-Scenario Cost Model

Year	Est. Total Completions (US E+SMB)	Base Case Cost/Resolution	Accelerated Cost/Resolution	Plateau Cost/Resolution	1-to-4 Rule Price Floor (Base)
2026	~117B	$0.25	$0.25	$0.25	$1.00
2027	~210B	$0.19	$0.13	$0.23	$0.76
2028	~410B	$0.16	$0.08	$0.22	$0.64
2029	~720B	$0.14	$0.06	$0.21	$0.56
2030	~1.1–4.2T	$0.12	$0.04	$0.20	$0.48

The margin discipline implication is consistent across all three scenarios. The 1-to-4 Rule, which requires that revenue per resolution equals or exceeds four times the AI cost to serve that resolution to maintain 75% gross margins, remains structurally sound in both the base and accelerated scenarios. In the plateau scenario, vendors who have already committed to $0.99 per resolution pricing face sustained margin pressure as token consumption growth prevents the cost-to-serve improvements their pricing assumed.

Section 6: The Governance Gap

Volume without governance is not an asset. It is a liability. As the volume of agentic task completions scales from billions to trillions, the absence of definitional infrastructure does not become more manageable. It becomes more dangerous. The governance gap in the resolution economy has four direct commercial costs that are already materializing in enterprise deployments.

41%	33%	12–19pp	40%+
DISPUTED COMPLETIONS	MISATTRIBUTED OUTCOMES	PILOT-TO-PRODUCTION LOSS	PROJECTS AT CANCELLATION RISK
Forrester root-cause analysis: 41% of failed agentic AI deployments are attributable to unclear success criteria. The absence of a completion definition makes every resolution a potential dispute.	33% of failed deployments trace to insufficient tool or data access, which produces misattributed completions: outcomes that were not actually agent-driven but are recorded as such.	Gartner cohort data: programs achieving 80%+ pilot accuracy lose 12 to 19 percentage points on launch to broader user populations. The governance infrastructure built for the pilot does not scale to production.	Gartner projects over 40% of agentic AI projects will be cancelled by end of 2027. The primary cause is not model quality. It is governance failures: absent audit trails, undefined permissions, and no explainability for high-stakes agent decisions.

The governance gap is not a future problem. Enterprises that launched agentic pilots in 2025 without robust audit trail infrastructure are now discovering in 2026 that the path to broader deployment requires rebuilding the permission and logging architecture they skipped in the rush to demonstrate capability. The specific gaps blocking scale in H1 2026 include audit trails that cannot satisfy enterprise procurement review, permission enforcement that cannot document role-based action authorization at the task level, and explainability frameworks that cannot answer why a specific agent took a specific action on a specific record.7

The SMB governance picture is more acute. Only 27% of small businesses report confidence in effective AI adoption, a figure that correlates directly with the absence of internal technical infrastructure to define, measure, or audit agent actions. An SMB deploying Intercom Fin or Zendesk AI on a per-resolution billing model is exposed to billing disputes it cannot resolve and cost variance it cannot predict, because it has no framework for distinguishing what the agent completed from what the agent attempted.

The Definitional Fragmentation Problem at Scale

The market currently has at least five competing measurement frameworks operating simultaneously: Salesforce's AWU, Intercom's per-resolution event, Zendesk's automated resolution metric, Microsoft's assisted action count, and ServiceNow's flow execution measure. Each is proprietary. None are interoperable. An enterprise running AI agents across multiple platforms cannot produce a consolidated view of its agentic activity volume, cost, or outcome quality. It cannot answer the most basic governance question: how many problems did our AI solve today, and what did it cost to solve each one?

> ““
>
> *Resolution as a Service is not a pricing model layered on top of existing agentic deployments. It is the architectural framework that makes agentic deployments governable, auditable, and commercially durable. The distinction matters because pricing model changes without architectural transformation do not solve the governance gap. They create a new billing surface on top of the same definitional chaos.*

Section 7: Vendor Adoption Tier Taxonomy

The vendor landscape for outcome-based AI pricing is stratified into three distinct tiers reflecting different degrees of commitment to resolution-based commercial models. The taxonomy is based on publicly disclosed pricing structures, earnings disclosures, and announced product strategies as of May 2026.

01

Billing Per Resolution Today

ACTIVE

Intercom · Zendesk · Forethought · Freshdesk (AI Copilot add-on)

These vendors have committed to published, per-resolution pricing for AI agent task completions. Intercom charges $0.99 per resolved conversation with a 50-resolution monthly minimum. Zendesk charges $1.50 to $2.00 per automated resolution. Both have explicit completion criteria, though neither fully satisfies all three Verifiable-Attributable-Finite criteria. Intercom's shift to outcome-based pricing for Fin drove net revenue retention from 112% to 146%, the strongest commercial validation of the model available from disclosed data.8 The limitation: these are all customer support use cases. Outcome-based pricing has not yet migrated to legal, finance, HR, or enterprise operations at Tier 1 scale.

02

Hybrid Models, Still Exposed

TRANSITIONING

Salesforce · ServiceNow · Microsoft · HubSpot · Workday

These vendors have layered outcome-adjacent pricing onto existing seat-based models without fully committing to resolution as the primary commercial unit. Salesforce's Agentforce offers both consumption-based and seat-based SKUs, with 9,500 paid deals as of Q3 FY26, growing 50% quarter-over-quarter. ServiceNow reports 50% of net new business is already non-seat-based. Microsoft Copilot is priced per user per month, not per task completion. Vendors in this tier retain significant Business Model Debt while building outcome-capable infrastructure. The January and February 2026 SaaSpocalypse disproportionately affected vendors in this tier, as capital markets began pricing the seat compression risk their hybrid structures acknowledge but do not resolve.

03

Pure Seat, High Exposure

VULNERABLE

Legacy ERP vendors · Most mid-market SaaS · Point solutions without AI-native architecture

The majority of enterprise SaaS vendors by count remain on pure seat-based pricing with no published outcome-based alternative. The psychological profile is avoidance: the hope that the transition moves slowly enough that it does not require action now. The data does not support that hope. Gartner projects 70% of vendors will refactor pricing away from pure per-seat models by 2028. IDC forecasts that seat-based pricing will be effectively dead for most software categories by 2028 to 2029. The 2029 renewal cycle is the structural forcing event: companies still defending pure per-seat models at that renewal will face accelerated churn from customers who have deployed AI agents and can document the resulting headcount reduction.

Conclusion: The Case for Resolution Discipline

The unmanaged resolution economy is not a future state. It is the present condition of enterprise and SMB AI deployment in 2026. Approximately 117 billion AI agent task completions are estimated to occur across US enterprise and SMB markets this year. Fewer than 3% carry any outcome-based pricing. The remaining 97% are being absorbed into contract structures that were designed for human-paced workflows and have no mechanism for capturing or auditing agentic value delivery.

The trajectory through 2030 makes the governance imperative urgent, not optional. Total completion volume is projected to reach between 1.1 trillion and 4.2 trillion annually by 2030. Inference costs will decline substantially across all three scenarios modeled in this report. But falling unit costs do not automatically produce governance clarity. They produce higher volumes of ungoverned completions at lower cost per event — a worse version of the same problem at a larger scale.

Resolution discipline is not a philosophical preference. It is the structural requirement for converting agentic AI deployment into durable, auditable, priceable business value. The market needs a framework in which every AI agent task completion can be asked three questions: Was it verifiable? Was it attributable? Was it finite? Until those three questions have defensible answers, the resolution economy will remain unmanaged. And an unmanaged economy of trillions of automated completions is not a foundation for enterprise trust.

The Path to Resolution Discipline: 2026–2030

2026	**Definitional Chaos at Scale** ~117B completions annually. Five competing proprietary measurement frameworks. Fewer than 3% priced as resolutions. 40% of projects at cancellation risk due to governance failures. The governance gap is documented; the solution is not yet deployed.
2027–28	**The Forcing Events** Gartner projects 40%+ of agentic projects cancelled. The 2028–2029 renewal cycle surfaces seat compression in contracts signed before AI agent deployment. Vendors without a resolution-based alternative face accelerated churn. Enterprise procurement teams begin requiring audit trail documentation as a condition of AI agent approval.
2029–30	**Resolution as the Default Commercial Unit** IDC projects seat-based pricing effectively dead for most software categories by 2028–2029. Bloomberg directional data: outcome-based pricing shifts from ~10% to ~60% of software pricing models by 2035. Vendors who completed the transition with defined resolution architectures compound. Those who deferred face the choice between distressed restructuring and exit.

From Measurement to Commercial Readiness

Identifying the governance gap is the diagnostic half of the problem. The operational half is moving from untracked completions to a defined resolution stack that can anchor a contract, survive a procurement review, and hold up in a billing dispute.

The output of a properly structured resolution diagnostic is a prioritized stack divided into three categories. First: outcomes that satisfy all three Atomic Resolution criteria immediately and are ready for contract language. Second: outcomes that are close but require additional instrumentation or quality validation before they can be priced per event. Third: outcomes that represent genuine platform value but cannot meet the verifiability and attribution standard required for per-resolution billing — these belong in a base platform fee, properly documented.

The most common failure mode in outcome-based pricing transitions is the conflation of all three categories into a single undifferentiated billing unit. Vendors who price outcomes from category two or three as if they were category one produce billing disputes, customer distrust, and the kind of governance failures that Gartner documents as the primary cause of agentic project cancellations. The diagnostic is not primarily a commercial exercise. It is a governance exercise with commercial consequences.

FOR VENDORS AND FOUNDERS

The resolution diagnostic described here is the entry point to the Resolution as a Service transition for seat-based vendors and the architectural foundation for founders building outcome-based pricing from the start. CPAG's advisory framework includes structured tooling for completing this diagnostic and translating its output into commercial and legal-review-ready documentation. Access is available at crownpointadvisorygroup.com.

SOURCES AND NOTES

1 Salesforce Q4 FY26 Earnings Release, February 25, 2026. AWU defined as "discrete tasks executed by AI agents in production across the Salesforce platform, including Agentforce and Slack."

2 Gartner press release, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026," August 26, 2025. McKinsey Global AI Survey 2026. G2 Enterprise AI Agents Report, October–November 2025.

3 Intercom disclosed 40M+ resolved conversations as of December 2025. Fin ARR surpassed $100M growing 350% year-over-year. Source: Sacra Intercom equity research, April 27, 2026. Salesforce Fortune article, "AI's Next Act," April 2026. Salesforce Q1 FY26 earnings, May 28, 2025.

4 SBA Office of Advocacy Business Trends and Outlook Survey (BTOS), September 2025. business.com / Dialog 2026 Small Business AI Outlook Report. U.S. Chamber of Commerce Empowering Small Business Report 2025.

5 McKinsey Global AI Survey 2026 (6.4 median hours saved weekly). Forbes / SMB Group 2024. SBA Office of Advocacy, September 2025.

6 Gartner March 2026 inference cost forecast. Epoch AI benchmark price analysis, arXiv November 2025 / March 2026. Statista AI model pricing data. Andreessen Horowitz "LLMflation" post.

7 Enterprise AI Agents 2026: Mid-Year Report on What's Working, Ampcome, May 2026. Forrester Agentic AI Wave Q1 2026. Gartner cohort study on pilot-to-production accuracy loss, 2026.

8 Intercom outcome pricing NRR improvement: Sacra Intercom equity research, April 27, 2026. Zendesk per-resolution pricing: CorePiper Zendesk AI Agent Pricing 2026 guide. ServiceNow non-seat revenue and IDC/Bloomberg forecasts: CPAG Research Brief, The Seat Is Dead, April 2026.

Legal Notice and Strategic Disclosures

This research brief is produced by Crown Point Advisory Group LLC for informational and thought leadership purposes only. It does not constitute financial, legal, investment, tax, or professional management consulting advice of any kind. Readers should conduct their own due diligence and consult qualified advisors before making any business or investment decisions based on the content of this brief.

Market sizing figures, volume estimates, and cost projections presented in this brief are CPAG Research estimates constructed through the methodology described in Section 2. These estimates represent directional ranges, not audited totals. Where individual figures are sourced from third-party analysts or vendor disclosures, those sources are identified in the footnotes. Aggregate estimates derived from CPAG methodology should be treated as directional indicators and validated against your own operational or vertical data before formal application.

The vendor analysis contained in this brief reflects publicly available pricing structures, earnings disclosures, and product announcements as of May 2026. Vendor strategies, pricing models, and market positions may have changed subsequent to publication. Crown Point Advisory Group makes no representation that the vendor information contained herein is current as of any date after publication.

Engagement with this document, or the localized application of its frameworks, does not establish a formal consultant-client or fiduciary relationship between the reader and Crown Point Advisory Group LLC or its affiliates. Formal advisory engagements require a separate, signed Master Services Agreement and a defined Statement of Work.

Portions of the data synthesis, structural modeling, and illustrative examples within this research brief were developed with the assistance of generative AI systems under strict editorial control and strategic direction by CPAG. All factual claims and estimates are the editorial responsibility of Crown Point Advisory Group. CPAG Research estimates are identified as such throughout and should be treated as directional indicators rather than audited figures.

All content, terminology, and proprietary frameworks contained in this brief — including Resolution as a Service, Atomic Resolution, the 1-to-4 Rule, the High-Fidelity Repository, the Biological Middleware Tax, the Friction Economy, and the Unmanaged Resolution Economy — are the intellectual property of Crown Point Advisory Group LLC. No part of this brief may be reproduced or redistributed for commercial gain without express written consent.