

NEWMIND RESEARCH · BENCHMARK REPORT 03

CUSTOMER SERVICE **COST-TO-SERVE** BENCHMARK 2026

What does each resolved ticket cost you, and how much could AI resolve? Benchmarks, a cost model for three sizes, playbooks and tools.

Edition 1

Desk-research edition

October 2026

Published

Licence

Single company use

CONTENTS

What's inside

PART A • THE BENCHMARK

01	About this edition	5
02	Executive summary	7
03	The market	9
04	Definitions and method	11
05	Cost per contact: assisted vs self-service	13
06	Agent labour cost	15
07	Channel mix and preference	18
08	Containment and deflection	20
09	Handle time and productivity	22
10	CSAT, FCR and quality	24
11	AI in customer service: evidence and predictions	26

PART B • SECTOR DEEP DIVES

12	Sectors at a glance	28
13	Retail and e-commerce	29
14	SaaS and technology	31
15	Hospitality and travel	33
16	Financial services and insurance	35
17	Logistics and transport	37

PART C • COST OF THE GAP AND PLAYBOOKS

18	The cost of the gap: three sizes	39
19	Scenario grid	41
20	Playbook 1: Measure cost per resolved issue	43
21	Playbook 2: Give agents AI assistance first	44
22	Playbook 3: Self-service that resolves	45
23	Playbook 4: Treat knowledge as infrastructure	46
24	Playbook 5: Hand-off and the right to a human	47
25	Playbook 6: Staffing and peaks after AI	48
26	Playbook 7: Quality assurance for AI conversations	49
27	Vendor landscape	50
28	Pricing models: seats, sessions and outcomes	51
29	Regulation: the European Union	52
30	Regulation: the United States	54
31	Roadmap: 90 days, 6 months, 12 months	55

PART D • WORKED EXAMPLES, CALCULATORS AND METHODS

32	Worked example: small team (1,000 tickets a month)	56
33	Worked example: mid-size operation (5,000 tickets a month)	59
34	Worked example: large operation (20,000 tickets a month)	62
35	Calculator: cost per resolved issue	65
36	Calculator: containment break-even	66
37	Calculator: staffing and workload	67
38	Calculator: labour cost per contact	68
39	Calculator: interval staffing	70
40	Method: building your own cost per contact	72
41	Method: measuring true containment	74
42	Method: finding and costing repeat contacts	75
43	Self-assessment: cost-to-serve maturity	76
44	Board questions and short answers	77
45	Sector cards: cost per resolved issue and first moves	78
46	Template: vendor demo test script	79
47	Template: monthly cost-to-serve review	81

PART E • TOOLKIT, TEMPLATES AND APPENDIX

48	Toolkit: KPI sheet	82
49	Toolkit: AI readiness scoring worksheet	84
50	Toolkit: extended scenario tables	85
51	Toolkit: vendor scorecard	87
52	Toolkit: contact-reason triage worksheet	88
53	Toolkit: business case template	89
54	Template: vendor question bank	90
55	Template: AI conversation quality rubric	93
56	Checklist: AI disclosure and hand-off	94
57	Template: pilot design	95
58	Template: risk register	96
59	Template: twelve-month KPI tracker	97
60	Template: contact-reason taxonomy starter	99
61	Template: knowledge article and content audit	100
62	Checklist: data readiness	101
63	Template: change plan and responsibilities	102
64	Appendix: model formulas	103
65	Appendix: source register and exclusions	104

REFERENCE

66	Glossary	108
67	Sources	109

Figures from third-party research carry a numbered source (full list in the Sources chapter). Everything else is labelled as NEWMIND analysis or model.

About this edition

This is Edition 1 of the NEWMIND Customer Service Cost-to-Serve Benchmark. It is a desk-research edition: it brings together published cost benchmarks, official labour statistics, company filings, peer-reviewed research and regulation, and turns them into a cost model you can run on your own numbers.

What this edition contains

- **Cost benchmarks:** the median cost of an assisted and a self-service contact¹, first contact resolution² and what they mean for the cost of a resolved issue.
- **Labour cost:** wages and employer cost per hour for US customer service representatives, from official statistics^{3,4}.
- **AI evidence:** the largest field study of AI assistance in a contact centre⁵, customer and leader surveys, and forecasts clearly labelled as predictions.
- **5 sector deep dives:** Retail and e-commerce, SaaS and technology, Hospitality and travel, Financial services and insurance, Logistics and transport.
- **Cost of the gap:** a transparent model for three operation sizes, a scenario grid and seven playbooks.
- **Vendor landscape, regulation and roadmap:** pricing models, vendor claims labelled as such, EU and US rules, and a 90-day, 6-month and 12-month plan.
- **Worked examples and calculators:** three operations walked through step by step with a month-by-month first year, and lookup tables for cost per resolved issue, containment break-even, staffing and labour cost per contact.
- **Methods:** interval staffing with Erlang C, building your own cost per contact, measuring true containment, costing repeat contacts, a maturity self-assessment and board questions.
- **Toolkit and templates:** a KPI sheet and tracker, a readiness scoring worksheet, scenario tables, a vendor scorecard and question bank, a quality rubric, a disclosure checklist, a pilot design, a risk register, a triage worksheet and taxonomy, content and data checklists, a change plan and a business case template.

What it does not contain yet

No original survey data. Edition 1 is desk research only and does not include our own survey. Edition 2 is planned to add a survey of service and operations leaders, with measures such as cost per resolved issue and true containment that companies do not publish.

How to read the numbers

- **Verified at the source.** The report draws on 41 sources. Every figure from its 36 published sources was checked against a saved copy of the source page; the other 5 are official statistics and company filings read from saved data captures. Figures we could not verify are listed as excluded and never used.
- **Labels.** Forecasts are marked as predictions. Numbers published by vendors and companies about themselves are marked as vendor or company claims.

- **No false trends.** Figures from different studies are never joined into a trend, and estimates from different firms are shown side by side, never averaged.
- **Models** are transparent: every formula and assumption is printed so you can replace it with your own data.

Who it is for

COOs, heads of operations and customer service leaders running 1,000 or more tickets a month, in retail and e-commerce, SaaS and technology, hospitality, financial services and insurance, and logistics and transport.

Licence: single company use. Please cite as: NEWMIND Research, Customer Service Cost-to-Serve Benchmark 2026, Edition 1.

Executive summary

Every customer contact has a price, and every repeat contact pays it twice. Cost per resolved issue is the number that turns service performance into money. That makes it the clearest place to benchmark, and to test how much AI can really resolve.

USD 13.50

Median cost of one assisted contact¹

USD 19.01

Cost of one resolved issue at average FCR (NEWMIND model)

14%

More issues resolved per hour with AI assistance⁵

Ten findings

01 An assisted contact costs more than seven times a self-service one

Gartner's benchmark puts the median cost at USD 13.50 per assisted contact and USD 1.84 per self-service contact¹. NEWMIND analysis: the gap is the single largest lever in the cost-to-serve.

02 The real unit is the resolved issue, not the ticket

At the industry-average first contact resolution of 71%², one resolved issue costs about USD 19.01 in assisted contacts (NEWMIND model). World-class operations reach 80%².

03 Fewer than one in ten customers finish in self-service

Only 9% of customers in one Gartner survey said they solved their issue completely in self-service⁶. A self-service attempt that fails costs both channels.

04 Customers are willing; far fewer actually use the bot

In a Gartner survey, 49% of customers said they would have been willing to use a chatbot, yet only 7% used one⁷.

05 AI assistance makes agents measurably faster, new agents most

In a study of 5,179 agents, AI assistance raised issues resolved per hour by 14% on average and by 34% for novice and low-skilled workers⁵.

06 Labour is the cost base AI acts on

The median US customer service representative earns USD 21.53 an hour³; with benefits, NEWMIND estimates about USD 31.48 an hour of employer cost⁴.

07 AI spend is outgrowing service budgets

Service leaders expect AI spending to grow 38% against 2% for their overall budget⁸. Only 24% demonstrated positive financial returns⁹.

08 Headcount cuts are the exception, not the rule

Just 20% of organisations reported reduced agent headcount due to AI¹⁰. NEWMIND analysis: savings appear when staffing plans change, not when a tool goes live.

09 Customers still want a way to reach a person

87% of customers say companies using generative AI in service must provide access to a human agent¹¹; 54% trust human agents more than AI, against 32% for AI¹².

10 Rules on AI in service now have dates

The EU AI Act applies generally from 2 August 2026, including the duty to tell people they are interacting with an AI system¹³. In the US, the FTC's Operation AI Comply targets deceptive AI claims¹⁴.

What to do on Monday

1. Compute your cost per contact and per resolved issue the way this report does (chapter 5).
2. Measure true containment, not deflection: count repeat contacts within seven days (chapter 8).
3. Price your gap for your volume with the three-size model (chapter 18).
4. Check your AI disclosure and route to a person in every flow against EU and US rules (chapters 29 and 30).
5. Set a baseline you will report against every month (chapter 48).

The market

Customer service is one of the largest labour pools in any economy and one of the first functions to be rebuilt around AI. In the United States alone, the Bureau of Labor Statistics counts 2.60 million wage-and-salary customer service representatives, with a median annual wage of USD 44,770³. Gartner's cost benchmark puts the median customer service and support budget at 0.7% of company revenue¹: small as a share, large in absolute terms, and almost entirely driven by contact volume.

USD 13.50

median cost of one assisted contact

USD 1.84

median cost of one self-service contact

0.7%

of revenue spent on service and support (median)

Sources for the three figures above: Gartner, Benchmarks to Assess Your Customer Service Costs¹.

AI for customer service: two published size estimates

Two research firms publish a public summary of the market for AI software used in customer service. They define the market differently and use different horizons, so we show both and never average them.

AI for customer service: 2024 market size by research firm (USD billion)



FIRM	2024 ESTIMATE	HORIZON FORECAST (FIRM'S PROJECTION)	GROWTH RATE (CAGR)
MarketsandMarkets ¹⁵	USD 12.1B	USD 47.8B by 2030	25.8%
Grand View Research ¹⁶	USD 13.0B	USD 83.9B by 2033	23.2%

Both estimates imply growth well above the 2% rise in overall service budgets that Gartner reports for the same function (see below). That gap is the story of this report: money is moving from labour to software inside budgets that are barely growing.

Outsourcers and contact-centre software: revenue from filings

Three US-listed companies give a public, audited view of the market around customer service: two customer experience outsourcers and one contact-centre software company. Figures are annual revenue from contracts with customers, as filed with the SEC. Concentrix's fiscal year ends on 30 November, so its years are not calendar years; the column headings follow the SEC calendar-year frame.


COMPANY	ACTIVITY	CY2022	CY2023	CY2024	CY2025
Concentrix Corporation ¹⁷	customer experience outsourcing	USD 6.32B	USD 7.11B	USD 9.62B	USD 9.83B
TTEC Holdings ¹⁸	customer experience outsourcing	USD 2.44B	USD 2.46B	USD 2.21B	USD 2.14B
Five9 ¹⁹	contact-centre software	USD 779M	USD 910M	USD 1.04B	USD 1.15B

Concentrix revenue rose in every year shown, from USD 6.32B to USD 9.83B. TTEC revenue was lower in the latest year (USD 2.14B) than in the first (USD 2.44B). Five9 grew in every year, from USD 779M to USD 1.15B. We report the direction only; filings explain the drivers, and we do not attribute the changes to AI.

Adoption is broad and still rising

Stanford's AI Index reports that 78% of organisations used AI in 2024, up from 55% the year before²⁰. Official EU statistics are more conservative because they count enterprises of every size: 20.0% of EU enterprises used AI technologies in 2025, an increase of 6.47 percentage points on the year before²¹. Size matters: large enterprises are more than three times as likely as small ones to use AI.

Share of EU enterprises using AI technologies in 2025, by size (%)



Budgets: AI spend is growing much faster than the function

- AI spending by customer service leaders rose by 38%, while overall service and support budgets rose by just 2% (survey of 199 leaders)⁸.
- Service and support leaders invested a median of 12% of their 2025 budget in AI, the highest of the ten business functions Gartner assessed⁹.
- Only 24% of service and support leaders demonstrated positive financial returns across their AI use cases⁹.
- **Prediction:** Gartner predicts that over 50% of customer service organisations will double their technology spend by 2028¹⁰.

The cost of the technology is falling fast

The AI Index reports that the inference cost for a system performing at the level of GPT-3.5 dropped over 280-fold between November 2022 and October 2024²⁰. For a support leader, this means the per-conversation cost of the model itself is rarely the constraint any more. The constraint is the quality of the knowledge the model draws on, the integrations that let it act, and the hand-off to a human.

What this means for you

Your budget will probably not grow much. The savings have to come from inside it: fewer assisted contacts, faster resolution of the ones that remain, and fewer repeat contacts. The rest of Part A measures each of these.

Definitions and method

Customer service metrics are easy to quote and hard to compare. The same word, “deflection” for example, can mean four different things depending on who measures it. This chapter fixes the definitions used in the rest of the report and explains how every figure was sourced.

Definitions used in this report

TERM	MEANING IN THIS REPORT
Contact	One customer interaction in any channel (call, chat, email, message, self-service session).
Ticket	A customer issue logged in a help desk. This report treats one ticket as one contact unless stated; repeat contacts are handled through first contact resolution.
Assisted contact	A contact handled by a human agent (phone, chat, email, social).
Self-service contact	A contact handled without a human: help centre, chatbot, AI agent, app or portal.
Containment	Share of self-service sessions that end without a human contact in the same session.
Deflection	Share of would-be assisted contacts avoided because self-service answered first. Vendors often use it loosely.
Resolution	The issue is fixed and the customer does not need to come back about it.
First contact resolution (FCR)	Share of issues resolved in the first contact, with no repeat within a fixed window.
Average handle time (AHT)	Talk or chat time plus hold and after-contact work, per contact.
Resolutions per hour	Issues an agent resolves in one working hour; captures concurrent chats, which AHT does not.
Cost per contact	Total service cost divided by contacts, by channel type.
Fully loaded hour	Wage plus employer-paid benefits per working hour.
Cost per resolved ticket	Cost per contact divided by the FCR rate: the expected cost of one fixed issue.

How the figures were collected

- **Primary sources first.** Official statistics (BLS, Eurostat, US Census Bureau), filings (SEC EDGAR), regulators (EUR-Lex, FTC, CFPB), peer-reviewed or working-paper research (NBER), then analyst press releases and public summaries.
- **Exact excerpts.** Every number was checked against the text of its source page and stored with the page address, publisher, date and the excerpt. A number we could not read on the source page is not printed.
- **Labels.** Measured figures, survey results, predictions, company claims and vendor claims are labelled as such. Predictions are never presented as results; vendor claims are never presented as benchmarks.
- **No averaging across studies.** Figures from different studies, years or definitions are shown side by side.

- **Models are labelled.** Every calculated figure carries the label “Market research model · based on published benchmarks” and states its inputs.
- **USD only.** All amounts are in US dollars as published; nothing is converted.
- **Loaded hour.** The fully loaded hour is the BLS median hourly wage for customer service representatives (USD 21.53, 2025 annual)³ divided by the wages-and-salaries share of total compensation for private industry workers in the BLS employer cost survey (68.4%, Q2 2026)⁴: USD 31.48. It adds benefits at the all-private-industry ratio, not this occupation's; supervision, tools, facilities and overhead are not included.

Two counts of the same workforce

The BLS occupational wage survey counts 2,595,750 customer service representatives in its 2025 annual estimate³. O*NET OnLine shows 2,666,000 for 2025²². The two come from different programmes and definitions: the wage survey counts wage-and-salary jobs and excludes the self-employed. We use the wage survey for every wage figure and O*NET only for its outlook, which it rates as a decline of 1% or more over 2025 to 2035, with 289,500 projected job openings over the projection period²². Openings remain high in a shrinking occupation because of turnover.

Reading the NBER study correctly

The strongest causal evidence on AI in support comes from a study of 5,179 customer support agents at one company⁵. It is one firm, technical support for small businesses, with 89% of agents located outside the United States. Its effect sizes are evidence of what AI assistance can do for agents; measure your own team before you plan on them.

What Edition 1 does not include

- No original survey of support teams. Edition 2 will add one.
- No paid analyst reports: we use only the public summaries and press releases that firms publish.
- No figures from pages that refused automated access or could not be read; these are listed in the source register in Part D.

Cost per contact: assisted vs self-service

The single most useful number in customer service economics is the price gap between a contact handled by a person and one handled by self-service. Gartner's benchmark puts the median at USD 13.50 for assisted channels and USD 1.84 for self-service¹: an assisted contact costs about 7.3 times as much. Gartner also notes that assisted channels such as phone, chat and email tend to have similar costs per contact.

Median cost per contact, assisted vs self-service (USD)

Assisted channels



Do not trend the older figures

An earlier Gartner survey quoted an average of USD 8.01 per live contact and about USD 0.10 per self-service contact⁶. It is a different study with a different method, so the two pairs are not a time series. We use the current benchmark everywhere and show the older one only so that readers who have seen it know where it fits.

STUDY	ASSISTED / LIVE	SELF-SERVICE	USE
Gartner cost benchmark (1 February 2024) ¹	USD 13.50	USD 1.84	Current benchmark used in every model
Gartner customer survey (25 September 2019) ⁶	USD 8.01	USD 0.10	Older study, different method; shown for context only

The hidden cost: self-service that does not resolve

Self-service is only cheap when it finishes the job. In Gartner's customer survey only 9% of customers reported solving their issue completely via self-service⁶. If every failed attempt then becomes an assisted contact, the expected cost of sending a customer to self-service first is USD 1.84 plus 91% of USD 13.50, or USD 14.12. That is more than going straight to an agent.

Break-even point

At the median costs, a self-service attempt pays for itself once it fully resolves more than **13.6%** of the contacts it receives ($\text{USD } 1.84 \div \text{USD } 13.50$). Below that, it adds cost. Above it, every extra point of true resolution saves money.

Market research model · based on published benchmarks

What one resolved ticket costs

A contact is not a resolution. When the first contact fails, the customer comes back, and you pay again. If each assisted contact resolves the issue with a probability equal to the first contact resolution rate, the expected cost

of one resolved issue is the cost per contact divided by that rate. SQM Group, a vendor of contact-centre benchmarking services, reports an industry average FCR of 71% and treats 80% or higher as world-class². At the median assisted cost, that is USD 19.01 per resolved issue at average FCR and USD 16.88 at world-class FCR.

Cost per resolved issue at assisted cost, by first contact resolution rate (USD)



Market research model · based on published benchmarks. Inputs: Gartner median assisted cost per contact¹; SQM FCR rates². Assumption: each contact resolves independently with probability equal to FCR.

How to measure your own cost per contact


- Take the full annual cost of the service function: salaries, benefits, supervision, workforce management, training, recruitment, telephony and software, outsourcing fees and a share of facilities.
- Split it into assisted and self-service pools. Self-service carries its platform, content and AI usage costs.
- Divide each pool by its contacts for the same year. Use contacts, not tickets, so repeat contacts are counted.
- Divide assisted cost per contact by your FCR rate to get your cost per resolved ticket.
- Compare with the medians above, then use the KPI sheet in Part D to track it monthly.

A regulator's view of the same gap: the CFPB cites third-party reports estimating that chatbots save about USD 0.70 per customer interaction compared with human agents²³. This is a cited third-party estimate, not a CFPB finding.

Agent labour cost

Labour is most of the cost of an assisted contact. The BLS wage survey puts the median hourly wage of US customer service representatives at USD 21.53 and the median annual wage at USD 44,770 (2025 annual)³. The spread is narrow: the 90th percentile earns about twice the 10th.

Customer service representatives: annual wage distribution, United States, 2025 annual (USD)




MEASURE	HOURLY WAGE	ANNUAL WAGE
10th percentile	USD 15.27	USD 31,750
25th percentile	USD 17.72	USD 36,870
Median	USD 21.53	USD 44,770
Mean	USD 22.40	USD 46,590
75th percentile	USD 24.89	USD 51,780
90th percentile	USD 30.57	USD 63,590

Wage-and-salary employment; excludes self-employed. Industry rows are a selection relevant to the five report sectors.

Wages by industry

The same job pays differently by industry. Business support services, which includes contact centres and outsourcers, sits near the bottom; publishing (which includes software publishers) near the top.


Median annual wage of customer service representatives by industry (USD)



From wage to fully loaded hour

Employers pay benefits on top of wages. In the BLS employer cost survey for private industry, wages and salaries made up 68.4% of total compensation in Q2 2026, and benefits 31.6%⁴. Applying that ratio to the median customer service wage gives a fully loaded hour of about USD 31.48, or about USD 65,453 a year per agent before supervision, tools and facilities. The ratio is for all private industry workers, not this occupation alone.

Employer cost per hour worked, private industry, all workers (USD)



Employer cost per hour worked rose from USD 46.14 in Q1 2024 to USD 49.46 in Q2 2026⁴. Every year without a productivity gain makes each assisted contact more expensive.

Cross-check: what an hour of chat costs

In the NBER study, an agent handled a baseline of 2.6 chats per hour; with AI assistance this rose by 0.37⁵. At a US fully loaded hour of USD 31.48, that is about USD 12.11 of agent time per chat before AI and USD 10.60 with it. The first number lands close to Gartner's USD 13.50 median for an assisted contact¹, which suggests the benchmark is a sensible anchor for US teams.

Market research model · based on published benchmarks. Inputs: BLS median wage³; BLS wage share of compensation⁴; NBER chats per hour⁵. The NBER firm's agents were mostly outside the US, so its own costs were lower; the figures show what its productivity would mean at US wages.

Sources

Sources accessed 6 Oct 2026. Figures quoted are short, attributed and linked; no third-party tables or charts are reproduced.

1. Gartner, Benchmarks to Assess Your Customer Service Costs (1 Feb 2024). Research summary (public abstract). <https://www.gartner.com/en/documents/5164231>
2. SQM Group, First Call Resolution metric and operating philosophy (20 Aug 2026). Industry benchmark (vendor research). <https://www.sqmgroup.com/resources/library/blog/fcr-metric-operating-philosophy>
3. U.S. Bureau of Labor Statistics, Occupational Employment and Wage Statistics, Customer Service Representatives, 2025 annual (data accessed 6 Oct 2026). Official statistics. <https://www.bls.gov/oes/>
4. U.S. Bureau of Labor Statistics, Employer Costs for Employee Compensation, private industry workers, Q2 2026 (data accessed 6 Oct 2026). Official statistics. <https://www.bls.gov/news.release/ecec.toc.htm>
5. National Bureau of Economic Research (Brynjolfsson, Li and Raymond), Generative AI at Work, NBER Working Paper 31161 (1 Apr 2023). NBER working paper, not peer-reviewed. <https://www.nber.org/papers/w31161>
6. Gartner, Gartner Says Only 9% of Customers Report Solving Their Issues Completely via Self-Service (25 Sep 2019). Survey, press release. <https://www.gartner.com/en/newsroom/press-releases/2019-09-25-gartner-says-only-9-of-customers-report-solving-their>
7. Gartner, Gartner Finds Only 27% of Customers Would Try a Chatbot Again After a Negative Experience (2 Sep 2026). Survey, press release. <https://www.gartner.com/en/newsroom/press-releases/2026-09-02-gartner-finds-only-27-percent-of-customers-would-try-a-chatbot-a-gain-after-a-negative-experience>
8. Gartner, Gartner Survey Finds AI Spending by Customer Service Leaders Has Surged by 38% Despite Budgets Rising by Just 2% (26 Aug 2026). Survey, press release. <https://www.gartner.com/en/newsroom/press-releases/2026-08-26-gartner-survey-finds-ai-spending-by-customer-service-leaders-has-surged-by-38-percent-despite-overall-service-and-support-function-budgets-rising-by-just-2-percent>
9. Gartner, Gartner Survey Finds Customers Are Three Times More Likely to Use Third-Party GenAI Than Company-Provided Chatbots (8 Jul 2026). Survey, press release. <https://www.gartner.com/en/newsroom/press-releases/2026-07-08-gartner-survey-finds-customers-are-three-times-more-likely-to-use-third-party-genai-than-company-provided-chatbots-for-customer-service>
10. Gartner, Gartner Predicts Over 50% of Customer Service Organizations Will Double Their Technology Spend by 2028 (31 Mar 2026). Prediction and survey, press release. <https://www.gartner.com/en/newsroom/press-releases/2026-03-31-gartner-predicts-over-50-percent-of-customer-service-organizations-will-double-their-technology-spend-by-2028>
11. Gartner, Gartner Survey Finds 87% of Customers Say Companies Using GenAI for Customer Service Must Provide Access to a Human Agent (4 Aug 2026). Survey, press release. <https://www.gartner.com/en/newsroom/press-releases/2026-08-04-gartner-survey-finds-87-percent-of-customers-say-companies-using-genai-for-customer-service-must-provide-access-to-a-human-agent>
12. Gartner, Gartner Survey Finds 85% of Service and Support Leaders Are Expanding Human Agent Responsibilities (28 Apr 2026). Survey, press release. <https://www.gartner.com/en/newsroom/press-releases/2026-04-28-gartner-survey-finds-eighty-five-percent-of-service-and-support-leaders-are-expanding-human-agent-responsibilities-despite-expectations-of-mass-ai-layoffs>
13. European Union (EUR-Lex), Regulation (EU) 2024/1689 (Artificial Intelligence Act) (13 Jun 2024). Regulation (official text). <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
14. U.S. Federal Trade Commission, FTC Announces Crackdown on Deceptive AI Claims and Schemes (25 Sep 2024). Regulator press release. <https://www.ftc.gov/news-events/news/press-releases/2024/09/ftc-announces-crackdown-deceptive-ai-claims-schemes>
15. MarketsandMarkets, AI for Customer Service Market (1 Feb 2025). Market sizing (public summary). <https://www.marketsandmarkets.com/Market-Reports/ai-for-customer-service-market-244782574.html>
16. Grand View Research, AI for Customer Service Market Size Report (page accessed 6 Oct 2026). Market sizing (public summary). <https://www.grandviewresearch.com/industry-analysis/ai-customer-service-market-report>
17. U.S. Securities and Exchange Commission (EDGAR XBRL company facts), Concentrix Corporation, Revenue from contracts with customers (10-K) (data accessed 6 Oct 2026). Company filing. <https://www.sec.gov/edgar/search/>
18. U.S. Securities and Exchange Commission (EDGAR XBRL company facts), TTEC Holdings, Revenue from contracts with customers (10-K) (data accessed 6 Oct 2026). Company filing. <https://www.sec.gov/edgar/search/>
19. U.S. Securities and Exchange Commission (EDGAR XBRL company facts), Five9, Revenue from contracts with customers (10-K) (data accessed 6 Oct 2026). Company filing. <https://www.sec.gov/edgar/search/>
20. Stanford HAI, AI Index Report 2025 (1 Apr 2025). Academic index. <https://hai.stanford.edu/ai-index/2025-ai-index-report>
21. Eurostat, Use of artificial intelligence in enterprises (1 Dec 2025). Official statistics. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Use_of_artificial_intelligence_in_enterprises

22. O*NET OnLine (US Department of Labor), Customer Service Representatives (occupation profile) (page accessed 6 Oct 2026). Official occupational data. <https://www.onetonline.org/>
23. Consumer Financial Protection Bureau, Chatbots in consumer finance (issue spotlight) (6 Jun 2023). Regulator research. <https://www.consumerfinance.gov/data-research/research-reports/chatbots-in-consumer-finance/chatbots-in-consumer-finance/>

© 2026 NEWMIND PARTNERS. Licensed for use within one company. NEWMIND PARTNERS AG, Luzern, Switzerland · NEWMIND PARTNERS Inc., Sheridan, Wyoming, USA · hi@newmindpartners.com

Get the full report

This sample includes the cover, contents, about page, executive summary, the market, definitions and method, cost per contact, agent labour cost and the sources cited in them. The full report continues with channel mix, containment, handle time, quality, AI evidence and predictions, five sector deep dives, the cost of the gap for three sizes, a scenario grid, seven playbooks, the vendor landscape, EU and US regulation, a 12-month roadmap, worked examples for three sizes, four calculators and a ready-to-use toolkit of templates and checklists.

Full report

newmindpartners.com/benchmark-customer-service-cost.html#report

NEXT STEP

TURN THIS BENCHMARK INTO YOUR CUSTOM SERVICE WORKFLOW.

You now know what a contact costs and where the gap is. NEWMIND PARTNERS designs, builds and runs the workflow that closes it: agent assist, self-service that resolves, clean hand-offs and quality review, built around the systems you already use.

01

Feasibility audit

We measure your service operation on your data (read-only): cost per contact and per resolved issue, repeat contacts, true containment and handle time by reason. Then we compare every solution on the market, off-the-shelf or custom.

02

Design

We design the workflow that fits your help desk, your knowledge and your policies, with a visible route to a person and a human checkpoint on every decision that matters.

03

Build and run

Our senior engineers build it, connect it to the systems you already use, and run it in production with an audit trail.

04

Prove it

Baseline on day one, then the change in cost per resolved issue, hours and satisfaction, reported every month.

Start with a feasibility audit

Find out what each resolved ticket really costs you, which solution fits best, and the business case to back it. No commitment.

[Start a feasibility audit →](#)