


New 2026 Benchmark Maps AI Customer Service Performance Across Six Industries: Resolution, CSAT, and Cost Per Resolution

Aissist.io study across six industries and 40+ sources puts vendor headline resolution rates at 67–90% and estimates CSAT of 78

MOUNTAIN VIEW, CA, UNITED STATES, July 8, 2026 /EINPresswire.com/ -- San Jose, CA, UNITED STATES, [July 5th, 2026] -- Aissist.io, the agentic AI operational layer for customer service, today published its [AI Customer Service Benchmark](#) 2026, an industry-wide analysis of the three metrics that determine whether AI support actually works: genuine end-to-end resolution rate, customer satisfaction (CSAT), and cost per resolution.



Driving factors of AI resolution by industries

Synthesizing more than 40 sources — vendor disclosures, independent aggregates, contact-center cost studies, and named customer deployments published between 2024 and 2026 — the benchmark holds vendor marketing claims side-by-side against independently verified results, and discloses the limits of the data on both sides.

“

Resolution is the only number that pays the bill, and it deserves honest measurement. We published this benchmark with its caveats visible so buyers can compare systems instead of marketing claims”

Lifan

Key findings from the 2026 benchmark:

The claimed-versus-field gap is structural, not anomalous. Vendor headline rates of 67–90% come from high-structure case studies, while independent cross-program aggregates put the tier-1 automation median near 41%, with a top quartile around 59%. The report notes that

source definitions vary — some publishers report that same aggregate as deflection — and

instructs readers to treat it as a directional floor rather than an audited median. Vendor numbers are also moving: Intercom's Fin, which marketed a 67% average in 2025, now markets 76%.

Resolution is not deflection. Counting any conversation that never reached a human — including customers who gave up — inflates reported performance by 20 to 40 points versus counting only issues genuinely solved end-to-end. The benchmark tags or excludes deflection-based figures from its verified ranges.

INDUSTRY	RESOLUTION RATE	OVERALL CSAT	HUMAN \$/TICKET	AI \$/RESOLUTION
Ecommerce & Retail	70-84%	76-80	\$2.70-\$5.60	\$0.50-\$2
Consumer Fintech	60-75%	81-83	\$15-\$30	\$1-\$5
SaaS & Software	50-70%	78-80	\$18-\$35	\$1-\$3
Travel & Hospitality	45-70%	82/72	\$10-\$25	\$1-\$3
Telecom & Utilities	40-60%	62-68	\$20-\$30	\$1-\$3
Healthcare & Insurance	40-60%	57-81	\$20-\$40	\$1-\$4

Travel CSAT shown as hospitality / airlines. AI-handled CSAT runs 5-10 pts below overall baseline.

Benchmarks by industries

Industry rank-order follows intent structure and data transparency. Ecommerce and retail lead at 70-84% verified resolution because their highest-volume intents are structured and data-rich. Consumer fintech (60-75%) and SaaS (50-70%) follow, while telecom, utilities, healthcare, and insurance trail at 40-60%, where issues are ambiguous, regulated, or emotionally charged and the data needed to resolve them is fragmented.

Architecture moves the number as much as industry does. Agentic AI systems typically resolve 10-20 points higher than retrieval-based bots, multi-agent designs add another 10-15 points, and giving the AI the ability to take action — refunds, account updates, rescheduling — is worth 20-30 points of resolution versus information retrieval alone.

The honest cost of an AI resolution is about \$5, not the list price. Human-handled tickets range from roughly \$2.70 in retail to \$60 in complex B2B, while AI resolutions cost \$0.50-\$2.37 at the unit level. Counting connectors, engineering, and platform fees, the report estimates a realistic all-in figure near \$5 per AI resolution in many B2B deployments — clearly labeled as the report's own synthesis rather than a published benchmark, and still roughly six times cheaper than the ~\$30 human equivalent. A 2.3× repeat-contact rate means deflection that fails to resolve can quietly raise total cost while appearing cheaper on a dashboard.

AI-handled CSAT runs 5-10 points below human-handled for the same team, against a cross-industry average near 78/100. The fastest way to damage it is a weak escalation handoff that forces customers to re-explain their issue.

" 'Resolution is the only number that pays the bill, and it deserves honest measurement. We published this benchmark with its caveats visible — which figures are disclosed, which are reported, and which are our own estimates — so buyers can compare systems instead of marketing claims,' said Lifan, Co-Founder of Aissist.io."

The benchmark recommends that buyers pin a resolution definition — no customer reply within

a fixed window and no escalation — and run a 50-question evaluation set across their real top intents before comparing any vendor numbers. The full report, including per-industry data tables, a vendor claimed-versus-verified comparison, pricing model analysis, and complete methodology, is available at: <https://aissist.io/industries/ai-customer-service-benchmark-2026>

About Aissist.io

Aissist.io is the stack-agnostic agentic AI operational layer for customer service and sales, built to lift genuine end-to-end resolution across existing helpdesks rather than deflect tickets off a dashboard. Its platform wires three systems into a continuous improvement loop: AgentMesh (multi-agent automation), Pulse (deep insight and auto-diagnosis), and Evolve (continuous optimization), with native integrations across 10+ helpdesks including Intercom, Zendesk, Freshdesk, Kustomer, Gorgias, HubSpot, Salesforce, and Front. Aissist.io supports 65+ languages, offers outcome-based pricing, is ISO 27001 certified and GDPR compliant, and was named Best Agentic AI for Business by CIOReview in 2026. Learn more at <https://aissist.io>.

Aissist

Aissant Inc

[email us here](#)

This press release can be viewed online at: <https://www.einpresswire.com/article/924642663>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.