

# 52%

## FASTER TICKET RESOLUTION WITH AI

Industry benchmark across 10M+ support tickets analyzed

# The State of AI in Support Operations 2025 – 2026 Industry Report

AI productivity benchmarks · ROI data · analyst forecasts  
for IT help desk and customer service teams

## \$3.50

average ROI  
per \$1 invested in AI

## 80%

of routine issues resolved  
autonomously by AI by 2029  
(Gartner)

## 670 hrs

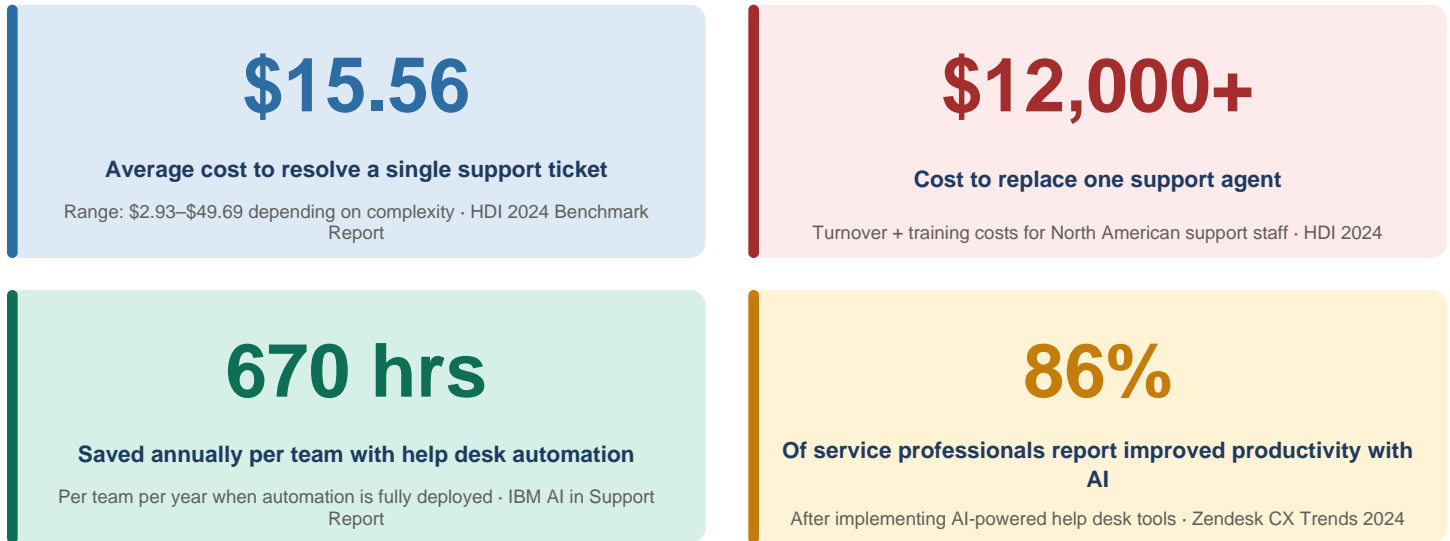
saved annually per  
team with automation  
(IBM)

### WHAT'S INSIDE

- The mounting pressure on support teams — and why you cannot hire your way out
- ROI data for each of the 5 core AI capabilities: triage, solutions, summaries, KB search, KB authoring
- Real-world benchmarks from McKinsey, Gartner, Freshworks, IBM, and Nielsen Norman Group
- Analyst forecasts: what the industry looks like in 2026, 2028, and 2029

Ticket volumes are rising. Customer expectations are compressing response-time windows. Yet support headcount isn't keeping pace — and the math of hiring your way out of the problem is getting harder to defend.

The data below establishes the baseline: what support operations cost today, where time goes, and why AI has moved from "nice to have" to a financial imperative for teams of every size.



## The growing gap: volume vs. capacity

- 52% of customers expect email responses within 1 hour
- Industry average actual response time: 12 hours, 10 minutes
- Trendsetting AI teams resolve tickets in 32 minutes avg.
- Non-AI teams: up to 36 hours for the same resolution
- 22% of all support tickets are now resolved with no human
- Over 40% of initial interactions handled by AI (Gartner)
- 93% of support leaders say automation will define the future
- 78% of customer service teams increased AI budgets in 2024

**"Trendsetting teams that use AI maintain response times under 20 seconds and resolve most queries in under 2 minutes."**

Freshworks CX Benchmark Report 2025

Research across millions of support tickets and thousands of agents identifies five AI capabilities with the clearest, most consistently measured impact on resolution time, agent productivity, and cost.

## 01 SUGGESTED SOLUTIONS FROM PAST TICKETS

AI surfaces relevant resolutions from ticket history before an agent opens a new case — instantly.

**+14%**

Agent issue resolution per hour

McKinsey study of 5,000 service agents

**+28%**

More conversations handled per agent

Google Cloud Agent Assist deployment data

**+43%**

Faster resolution with context-aware AI

vs. context-blind systems · Fini Labs 2024

**IMPACT IN PRACTICE:** Newer agents access institutional knowledge immediately — reducing escalations and onboarding time.

## 02 TICKET & CONVERSATION SUMMARIES

Long threads condensed instantly — agents pick up mid-conversation without reading everything.

**56%**

Of agents report time saved with AI summarization

Freshworks benchmark report

**-45%**

After-call work reduced by AI summarization tools

Dialpad deployment · retail sector

**-20%**

Handle time reduction from triage & summarization

Yielding \$150K annual labor savings

**IMPACT IN PRACTICE:** ServiceNow: 12–17 min of agent time saved per case. At 50 tickets/day, that's 10–14 hrs/agent/week.

## 03 INTELLIGENT TRIAGE & AUTO-CATEGORIZATION

AI detects intent, urgency, and category — routing tickets to the right queue in seconds, not minutes.

**95%**

AI category accuracy vs. 77% for human agents

Fini Labs benchmarks · Dec 2024

**35%**

Of tickets deflected by AI self-service tools

Zendesk CX Trends 2025 · McKinsey

**68%**

Deflection rate on employee requests at Equinix

Via AI-powered triage · Moveworks 2024

**IMPACT IN PRACTICE:** Manual triage: 3–5 min per ticket. AI triage: seconds. At 100 tickets/day, that's 5–8 hrs of coordinator time recovered.

## 04 AI-POWERED KNOWLEDGE BASE SEARCH

Natural language queries that understand intent — surfacing the right article, not keyword matches.

**-35%**

Support volume reduced with AI-powered KB search

Industry research · Pylon 2025

**-35%**

Time spent searching for information across the team

Knowledge management research

**+60%**

Higher ticket deflection with AI-first platforms

vs. traditional help desks · Gartner 2024

**IMPACT IN PRACTICE:** 91% of customers try self-service when it's well-organized. AI KB search is what makes self-service work at scale.

## 05 AI KNOWLEDGE BASE GAP ANALYSIS AND ARTICLE AUTHORIZING

Identifies missing Knowledge Base coverage for common issues, then generates draft articles from resolved ticket data.

**-60–80%**

Time to create a KB article with AI assistance

AI drafts need light review · Pylon 2025

**3–6 mo**

To reach "critical mass" (30%+ ticket deflection) with AI-built KB

Industry benchmark for KB maturity

**-23%**

Customer support tickets with a mature KB in place

Industry research across B2B platforms

**IMPACT IN PRACTICE:** Every resolved ticket becomes a reusable asset. Teams using AI authoring compound their KB value automatically over time.

### At a glance: combined impact of all five capabilities

Capability	Primary metric	Reported gain	Source
Suggested solutions	Issue resolution rate / hour	+14% (McKinsey)	
Conversation summaries	Handle time / after-call work	-20–45% (Freshworks / Dialpad)	
Intelligent triage	Ticket deflection rate	35–68% deflected (Zendesk / Moveworks)	
AI KB search	Support ticket volume	-35% (Pylon / Gartner)	
AI KB Gap Analysis & Authoring	Article creation time	-60–80% (Pylon 2025)	

Sources: McKinsey GenAI in Enterprise · Freshworks CX Benchmark · Dialpad case study · Zendesk CX Trends 2025 · Moveworks 2024 · Pylon 2025 · Gartner 2024

AI in support operations is no longer a speculative investment. Across thousands of documented deployments, the financial returns are measurable, repeatable, and often faster than expected.

**\$3.50**

**Average ROI per \$1 invested in AI**

With top performers achieving 8x returns

**3–6 mo**

**Typical payback period for AI deployment**

Most teams see positive ROI inside 6 months

**25–30%**

**Cost-to-serve reduction from AI implementation**

McKinsey: applying gen AI to customer service

## Real-world results

### Klarna

2.3M conversations in first month — equivalent to 700 FTE agents. Resolution time cut from 11 minutes to under 2 minutes. Estimated \$40M profit improvement cited in 2024 AI efficiency gains.

### Equinix

68% deflection on employee support requests. 43% autonomous resolution via AI-powered triage through Teams. Reported via Moveworks 2024 global deployment recap.

### Unity

\$1.3 million saved by deflecting 8,000 support tickets with AI. 90% employee satisfaction with IT. Frequently cited as benchmark for mid-to-large enterprise AI deployment ROI.

### NIB Health Insurance

\$22 million in savings — a 60% cost reduction — from AI-powered customer service transformation. One of the largest documented ROI outcomes in the customer service AI space.

## Estimate your own savings

**Annual Savings = (Cost per ticket × Monthly ticket volume × 12) × Reduction %**

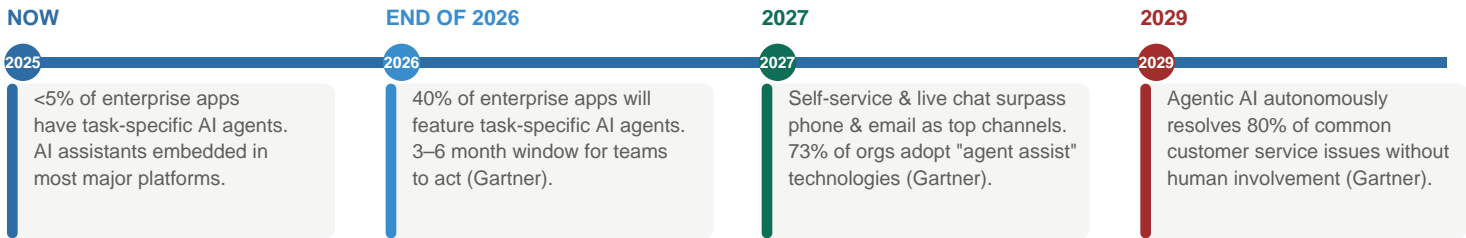
Using \$15.56 avg. ticket cost · 35% deflection rate: a team handling 2,000 tickets/month saves

**2,000 × \$15.56 × 12 × 0.35 = \$130,704 per year in support cost reduction**

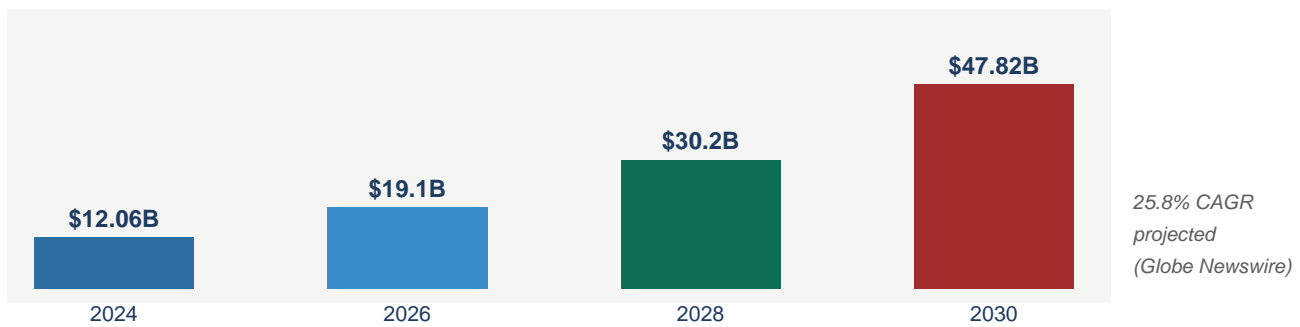
### Adoption context:

52% of businesses now evaluate helpdesk ROI within six months of deployment (HubSpot Research). 78% of customer service teams increased their helpdesk budgets in the past 12 months (ServiceNow Trends Report 2025).

The analyst consensus is unusually consistent: agentic AI in support is not a future possibility — it is a current deployment reality that is moving into mainstream enterprise use faster than most technology transitions in recent memory.



## Market growth: AI for customer service



### The window for "early mover" advantage is closing.

Teams deploying AI in 2025–2026 are building compounding advantages in resolution speed, KB quality, and customer satisfaction that will be difficult for slower-moving peers to close.



# Agentic AI for osTicket teams.

No migration. No per-agent pricing. No surprises.

## What Flexivity AI does

Flexivity AI brings the five AI capabilities documented in this report directly to your existing osTicket installation — for both IT help desk and customer-facing support — without requiring a platform migration.

All five capabilities are live in beta today:

- ✓ Suggested solutions from past tickets
- ✓ Ticket & conversation summaries
- ✓ Intelligent triage & auto-categorization
- ✓ AI-powered knowledge base search
- ✓ AI knowledge base gap analysis and article authoring

## A pricing model built differently

- **Predictable flat tiers**  
Pay for what the AI does across a capacity band — not per seat, not per resolution.
- **No spike risk**  
A product launch or seasonal surge does not change your bill. Costs tied to your team, not your customers' behaviour.
- **Scales with your team, not against it**  
Add agents freely. Your software cost stays constant as your headcount grows.

### Flexivity AI vs. migrating to a new platform

Migration for a mature support operation: \$10,000–\$150,000 in data migration, integrations, agent retraining, and productivity loss — before a single AI feature goes live.

**Flexivity AI adds AI to the platform your team already knows. Zero migration cost.**

## Join the beta — add agentic AI to your osTicket today.

Request early access or join our waitlist at

**[flexivity.ai/form/join-our-waiting-list](https://flexivity.ai/form/join-our-waiting-list)**

IT help desk · customer-facing support · both

This report cites third-party research for informational purposes. All statistics are attributed to their original sources. Flexivity AI makes no warranty as to accuracy of third-party data. © 2026 Flexivity AI. All rights reserved.