

Contact Center AI Benchmarks by Industry

2026.

Containment, AHT, CSAT-proxy and cost-per-contact benchmarks for AI-enabled contact centers, segmented by BFSI, insurance, healthcare, retail and BPO. Peer bands you can benchmark your own program against.

From strategy and pilots to production outcomes.

VP Contact Center

Head of CX

COO

CFO

6

Industries benchmarked

10 min

Read time · 39 charts

365 days

Post go-live tracking window

Two years of contact center AI hype produced a clearer picture than the vendor slides promised. Voice containment finally moved. Agent assist gains are real but bounded. And post-launch CSAT depends more on change management than model quality. This benchmark covers 600+ AI-enabled contact centers across BFSI, insurance, healthcare, retail and BPO.

6

Industries benchmarked

10 min

Read time · 39 charts

365 days

Post go-live tracking window

What this report answers

- AI containment rates by channel and industry
- AHT delta with and without agent assist
- CSAT-proxy movement post AI rollout
- Cost-per-contact bands by industry and estate size
- Where top-quartile programs are winning — and where they're stuck

01 Containment is real, but only for the right intents

Voice containment on transactional intents — balance, hours, appointment, status — now sits between 45% and 60% for well-designed programs. That is the first meaningful jump since the shift from touchtone to speech IVR.

Complex or empathetic intents still belong with humans. The programs winning at containment are the ones that segmented aggressively by intent, not the ones that chased a single global number.

02 Agent assist AHT gains are bounded

Median AHT reduction from agent assist is 15% to 25%. Programs that pushed past that ceiling usually did so by redesigning wrap-up or knowledge, not by adding more assist surfaces.

The failure mode we see most often is cognitive overload. Three assist panes is the ceiling before AHT starts moving the wrong way.

03 CSAT depends on change management, not the model

Post-launch CSAT movement varies by 20 points between similar deployments on the same platform. The differentiator is not the AI stack — it is how supervisors were trained, how coaching was rewired and how quickly the QA rubric was rewritten to reflect the new interaction pattern.

04 Cost per contact is a portfolio metric now

Blended cost per contact tells you less than it used to. AI containment shifts the mix toward complex, longer human interactions, which raises the human unit cost even as the total falls.

Report the metric three ways: contained, human-handled, blended. Boards that see only the blended number tend to under-invest in the human side of the estate.

- Contained AI interaction: usually 60–85% cheaper than a human-handled equivalent
- Human-handled interaction after containment lift: 5–15% more expensive per call
- Blended cost per contact: typically down 20–40% year over year in mature estates

05 The industry cuts that matter

BFSI leads on containment but lags on CSAT movement, mostly because compliance disclosures constrain conversational design.

Healthcare payers are the fastest-improving segment in 2026, driven by member-services automation. Retail leads on personalization; BPO leads on QA automation. Insurance is where the largest untapped ROI still sits.

THE BOTTOM LINE

The next twelve months of contact center AI will be won on operational discipline, not model selection. Rebuild your QA rubric first, then your assist stack, then your voice bot — in that order.

Questions we get asked

Can I benchmark my own program against this?

Yes — the appendix includes a self-assessment worksheet that maps your metrics to the bands.

Benchmark your own program. Model the business case with our ROI calculators, or book a working session with the practitioners who ship these programs.

pronix.ai/book
pronix.ai/roi

Pronix.ai is a specialized AI & CX Systems Integrator. Leading with AI. Innovating with Digital. Figures are drawn from delivered engagements and enterprise programs benchmarked by Pronix; ranges are reported instead of false precision. Third-party names and marks belong to their owners.