

State of AI in CX 2026

The money is moving. The returns are not. What the evidence says about why AI in customer experience underdelivers, and the six foundations that decide whether it pays back.



Data quality

Organisational
change

Process maturity

Use-case prioritisation

Channel architecture

Commercial
alignment

Prepared by

Rachel Lane, Founder and CEO

September 2026

Contents

03	Foreword
04	What we found
05	1. Data quality
06	2. Organisational change
07	3. Process maturity
08	4. Use-case prioritisation
09	5. Channel architecture
10	6. Commercial alignment
11	Practitioner tool: AI readiness scorecard
12	Practitioner tool: workflow verdict guide
13	Practitioner tool: talking to the board
14	Conclusion and 90-day roadmap
15	Methodology and references

Foreword

The foundation decides the return

I've been inside enough enterprise CX programmes to know what it looks like when AI deployment goes wrong. The frantic calls six weeks post go-live. The vendor blame game. The board presentation where nobody can explain why the numbers aren't moving.

2026 was meant to be the year that changed. Spending certainly did. Gartner found AI budgets in customer service rose 38% this year while the function's overall budget grew by 2%.² The money is moving. The returns are not keeping pace, and in most organisations nobody can say, in pounds, what the last round of AI investment returned.

This report pulls together the most credible published research of the past eighteen months and reads it the way a practitioner would. We organise it around the six dimensions we use to assess every programme: data quality, organisational change, process maturity, use-case prioritisation, channel architecture and commercial alignment.

The pattern is consistent. When AI in CX underdelivers, the cause most often sits beneath the model: in the data it reads, the process it runs, the people who own it and the numbers nobody agreed in advance. Every one of those is fixable, and every one of them is cheaper to fix before go-live than after.

Use the scorecard on page 11. Be honest with it. Then decide where your next pound goes.

ABOUT THIS REPORT

This is a synthesis of published research from Gartner, McKinsey, S&P Global, MIT Project NANDA and RAND, interpreted through the CXaiS six-dimension readiness framework. Every statistic is sourced on page 15. The interpretation, verdicts and tools are ours.

WHO IT IS FOR

Chief Customer Officers, COOs, CX and service leaders who own an AI budget, and the CX technology vendors who sell to them.



Rachel Lane

Founder and CEO, CXaiS

15 years designing and delivering enterprise CX technology programmes

What we found

Spending is up. Readiness is not.

AI has moved from experiment to expectation in customer experience. **91% of service leaders** report pressure from executive leadership to implement AI.¹ Yet the research is equally clear on the other side of the ledger. MIT Project NANDA found that **95% of integrated enterprise GenAI pilots** in its study had produced no measurable P&L impact,¹¹ and by some estimates more than **80% of AI projects** fail, around twice the rate of non-AI IT projects.¹²

Read together, the evidence points to six places where CX programmes lose their return. In most of them, the model is not the deciding factor.

- 1

63%

The clean extract illusion · Data quality
of organisations don't have, or aren't sure they have, the data management practices AI needs. Pilots run on prepared data. Production runs on the CRM.⁵
- 2

20%

The headcount mirage · Organisational change
of service leaders have actually reduced agent staffing because of AI. The business cases written in headcount are the ones most likely to be wrong.⁸
- 3

~3x

Automating the mess · Process maturity
High performers are nearly three times as likely to have redesigned workflows. AI accelerates whatever process it is given, broken handoffs included.¹⁰
- 4

42%

The pilot graveyard · Use-case prioritisation
of organisations reported abandoning most AI initiatives before production in 2025, up from 17% the year before. Use cases chosen by demo rarely survive contact with the P&L.⁹
- 5

3x

The front door has moved · Channel architecture
Customers are around three times more likely to start with third-party GenAI than a company chatbot. What reaches you is harder than it used to be.⁴
- 6

38% vs 2%

The budget squeeze · Commercial alignment
AI spend in customer service is up 38%. Function budgets are up 2%. The money is coming from somewhere, and only 39% of organisations see any EBIT impact.²

Across the programmes we assess, the most common problem is not a shortage of potential use cases. It is an inability to show which workflows are ready, what value is at stake and who owns the outcome.

THE CXAIS FIELD VIEW

DIMENSION 1 · DATA QUALITY

The clean extract illusion

The issue

An AI pilot almost always runs on a dataset someone prepared by hand. It performs well in the demo and the budget renews. Then production connects it to the systems the contact centre actually runs: a CRM that holds the same customer three different ways, a knowledge base last audited two reorganisations ago, and disposition codes agents choose to close the ticket rather than to describe it. The model performs exactly as built. The data underneath it was never ready.

63%

of organisations don't have, or are unsure they have, the right data management practices for AI⁵

60%

of AI projects unsupported by AI-ready data will be abandoned through 2026 (Gartner prediction)⁵

80%+

By some estimates, more than 80% of AI projects fail, around twice the rate of non-AI IT projects¹²

What good looks like

Data readiness is judged workflow by workflow, not for the organisation as a whole. The question stops being “is our data good?” and becomes “is the data this specific workflow reads accurate, current and complete enough for an AI to act on without a human checking it?” Knowledge content has a named owner and a review date. Contact reason codes describe what customers actually asked for.

The practitioner move

1. Map every data source each candidate workflow reads, including the knowledge articles.
2. Test a sample of live production records, never the prepared extract.
3. Grade each workflow's data GREEN, AMBER or RED before any build brief is written.

“The model will perform exactly as well as the knowledge article it retrieves. If nobody owns that article, nobody owns the outcome.”

THE CXAIS VIEW

● DIMENSION 2 · ORGANISATIONAL CHANGE

The headcount mirage

The issue

Many AI business cases in CX were written in headcount, because the staffing line is the easiest number to put in front of a CFO. It is also the number most likely to be wrong. When the people removed are the ones who knew where the process breaks, the savings disappear into repeat contact, complaints and escalations that nobody is left to handle.

20%

of service leaders have actually reduced agent staffing because of AI⁸

50%

of organisations that expected to significantly cut their service workforce will abandon those plans by 2027⁷

84%

of service leaders plan to add new skills to the agent role and change hiring profiles¹

What good looks like

Change is designed alongside the technology, not after it. Frontline roles are redesigned before go-live: who handles escalations, who tunes the AI, who owns the knowledge it relies on. Team leaders are measured on the new operating model rather than the old one. The business case counts capacity released and where it is redeployed, rather than roles removed.

The practitioner move

1. Name an accountable business owner for every AI workflow, outside IT.
2. Redesign the three frontline roles most affected before launch, including their KPIs.
3. Rebuild the business case around capacity and outcomes, then test it with Finance.

“If the business case only works when you cut the people who know where the process breaks, the business case doesn't work.”

THE CXAIS VIEW

DIMENSION 3 · PROCESS MATURITY

Automating the mess

The issue

AI accelerates whatever process it is given. If a service journey has five handoffs, three workarounds and a policy exception nobody documented, the AI will find all of them, at scale, in front of customers. Most organisations map the process as documented. The failures live in the process as it is actually run.

~3x

AI high performers are nearly three times as likely to have fundamentally redesigned workflows¹⁰

~6%

of organisations qualify as AI high performers, attributing 5% or more of EBIT to AI¹⁰

88%

of organisations now use AI in at least one business function¹⁰

What good looks like

Every workflow in scope has been mapped end to end as it really runs, including the handoffs between teams and systems where most failure hides. Exceptions are written down. Resolution is defined in the customer's terms, not the system's. This is the work we call workflow cartography and handoff forensics, and it happens before a single use case is built.

The practitioner move

1. Map the current state from live contacts and agent screens, not from the process manual.
2. Count the handoffs in each workflow and measure the failure rate at each one.
3. Fix or remove the worst handoff before automating around it.

“Most of the failures we are called in to diagnose six weeks after go-live were visible on the process map. Nobody drew it.”

THE CXAIS VIEW

● DIMENSION 4 · USE-CASE PRIORITISATION

The pilot graveyard

The issue

Use cases are too often chosen because a vendor demo impressed, a competitor made an announcement or the board asked a question. They are then judged against expectations that were never tied to the economics of a specific workflow. Agentic AI adds a new risk: Gartner describes widespread “agent washing”, where existing chatbots and automation are rebadged as agents.

42%

of organisations reported abandoning most AI initiatives before production in 2025, up from 17% in 2024⁹

46%

of AI proofs of concept are scrapped before production, on average⁹

40%+

of agentic AI projects will be cancelled by the end of 2027 (Gartner prediction)⁶

What good looks like

Use cases are ranked on three things together: the value at stake, the readiness of the workflow and the risk if the AI gets it wrong. The portfolio is sequenced in waves. Wave 1 contains only workflows that are GREEN on readiness and material on value. Everything else waits, with a named reason and a named fix.

The practitioner move

1. List every candidate workflow, not every idea. A use case without a workflow is not a use case.
2. Score each on value, readiness and risk, and record the evidence behind each score.
3. Commit only GREEN workflows to Wave 1, and write the build brief from that evidence.

“The discipline is saying no to the RED workflows, out loud, before anyone has spent money on them.”

THE CXAIS VIEW

DIMENSION 5 · CHANNEL ARCHITECTURE

The front door has moved

The issue

Customers now arrive having already asked ChatGPT, Gemini or Copilot. By the time they reach you, the easy question has been answered somewhere you don't control, and what remains is harder, more emotional or needs account access. Channel strategies built around deflecting simple contacts are optimising for a customer who increasingly isn't there.

~3x

Customers were around three times more likely to use third-party GenAI than a company chatbot in their latest service interaction⁴

87%

of customers say companies using GenAI must provide a route to a human agent³

74%

of B2B customers who use GenAI have used it to complete a task on their behalf (58% across all customers)³

What good looks like

Channels are designed as one system rather than a set of separate projects. AI captures intent and resolves only where confidence is high, with a visible route to a person and context that travels with the customer. Voice is treated as a priority AI channel. Public knowledge is structured so that third-party AI tools give customers correct answers about your products and policies.

The practitioner move

1. Map where customer journeys actually start, including outside your own channels.
2. Design the escalation path first, then the automation.
3. Measure repeat contact across all channels, not containment within one.

“Containment is the metric that ruined chatbots. If the customer calls back tomorrow, you didn't contain anything. You deferred it.”

THE CXAIS VIEW

DIMENSION 6 · COMMERCIAL ALIGNMENT

The budget squeeze

The issue

AI spending in customer service is rising far faster than the function's budget, which means the money is being taken from somewhere, usually labour and overhead. That raises the stakes on proving a return at exactly the moment most organisations can't. Activity metrics fill the gap: use cases launched, chats contained, handle time shaved. None of them is a number a CFO can bank.

38% vs
2%

AI spend by service leaders rose 38% while overall service budgets grew 2%²

39%

of organisations report any enterprise-level EBIT impact from AI¹⁰

95%

of integrated enterprise GenAI pilots in MIT NANDA's study had produced no measurable P&L impact¹¹

What good looks like

Value is quantified for each workflow before the build, baselined before go-live and tracked monthly after it. Finance agrees the measurement method in advance, so nobody argues about the number later. The CX leader can explain to the board, in pounds, what each AI workflow returned, and what was stopped and why.

The practitioner move

1. Establish a Value at Stake baseline for each workflow before build.
2. Agree the measurement method with Finance before go-live, not after.
3. Report monthly in P&L language. The translation table on page 13 is a starting point.

“Commercial accountability means the number is agreed before the build starts, not reverse-engineered after the board asks.”

THE CXAIS VIEW

AI readiness scorecard

Score your programme across the six dimensions

Score each statement from 0 (not in place) to 5 (fully embedded, with evidence you could show a sceptical CFO). Score what is true today, not what is planned. Add the twelve scores for a total out of 60. For a scored version with benchmarks and an indicative value at stake, take the free assessment at cxais.ai/free-assessment.

Statement	Score
 The data each AI workflow reads has been tested on live production records, not a prepared extract.	0 1 2 3 4 5
 Knowledge content used by AI has a named owner and a regular review cycle.	0 1 2 3 4 5
 Every AI workflow has an accountable business owner outside IT.	0 1 2 3 4 5
 Frontline roles, escalation paths and team leader KPIs have been redesigned for AI.	0 1 2 3 4 5
 In-scope workflows are mapped end to end as they actually run, including handoffs and exceptions.	0 1 2 3 4 5
 The worst-performing handoffs are fixed before automation, not after.	0 1 2 3 4 5
 Use cases are ranked on value, readiness and risk together, with evidence for each score.	0 1 2 3 4 5
 Wave 1 contains only workflows with a documented GREEN readiness verdict.	0 1 2 3 4 5
 AI and human channels share context, and customers can always reach a person.	0 1 2 3 4 5
 Success is measured on repeat contact across channels, not containment within one.	0 1 2 3 4 5
 Each AI workflow has a Value at Stake baseline agreed with Finance before build.	0 1 2 3 4 5
 AI returns are reported to the board in P&L terms at least quarterly.	0 1 2 3 4 5

Your maturity level

Level	Score	Profile	Next step
Explore	0 to 17	No formal AI strategy, governance or readiness assessment.	Run a baseline assessment before committing further spend.
Build	18 to 29	Foundations under way but uneven across dimensions.	Fix the lowest-scoring dimension before scaling any pilot.
Operational	30 to 41	AI live in production with some commercial discipline.	Introduce workflow verdicts and a Finance-agreed baseline.
Optimised	42 to 51	Measured, governed and improving.	Move reporting to monthly P&L terms and sequence Wave 2.
Leading	52 to 60	AI treated as a strategic competitive advantage.	Protect the lead: review models and data on a fixed cycle.

Workflow verdict guide

GREEN, AMBER or RED, before anything is built

A readiness verdict is given per workflow, not per programme. A workflow is only as ready as its weakest dimension, so one RED means the workflow waits, however strong the rest looks.

Dimension	GREEN	AMBER	RED
 Data quality	Live records tested; knowledge owned and current	Gaps known, with a fix scheduled	Untested, unowned or contradictory data
 Organisational change	Business owner named; roles redesigned	Owner named; roles not yet redesigned	No business owner, or IT-owned only
 Process maturity	Mapped as run; worst handoff fixed	Mapped; handoff failures unresolved	Not mapped, or mapped only as documented
 Use-case prioritisation	Material value, low risk, evidence recorded	Value clear, risk or readiness unproven	Chosen by demo, no value case
 Channel architecture	Human route and shared context designed	Human route exists, context lost on transfer	No route to a person, or containment as the goal
 Commercial alignment	Baseline agreed with Finance before build	Baseline estimated, not yet agreed	No baseline; value to be proven later

How to use the verdicts

All GREEN

Eligible for Wave 1. Write the build brief from the evidence gathered, including the baseline and the named owner.

Any AMBER

Wave 2 candidate. Assign each AMBER a fix, an owner and a date. Re-score when the fix lands.

Any RED

Not now. Record the reason. Stopping a RED workflow early is a result worth reporting to the board.

Talking to the board

Translate activity into P&L language

Boards fund returns, not activity. Only 39% of organisations can point to any EBIT impact from AI,¹⁰ and most CX reporting still leads with operational metrics that Finance cannot bank. Use this table to reframe your next update. Replace the placeholders with your own baselined numbers, and only with numbers Finance has agreed.

Metric	How it is usually reported	How to report it to the board
Containment	"Our AI assistant contained 40% of chats this quarter."	"AI resolved x contacts end to end with repeat contact held below y%, releasing z hours of capacity, now redeployed to retention calls."
Handle time	"Agent assist cut average handle time by 30 seconds."	"Agent assist is worth £x a year against the baseline Finance signed off in <i>month</i> , with first contact resolution up y points."
Deployment	"We launched twelve AI use cases this year."	"Four workflows are live and returning £x. Two were stopped at a RED verdict before build, avoiding an estimated £y of spend."
Satisfaction	"CSAT on the bot is 4.1 out of 5."	"Customers who used the AI route show x% lower churn intent than the matched group, worth £y in retained revenue."
Workforce	"AI will let us reduce headcount by 15%."	"AI absorbed x% volume growth with no new hires, and y agents moved into escalation and AI-tuning roles."

Every number in the right-hand column needs a baseline taken before go-live. If you don't have one, that is the first thing to fix, and it is still fixable for your next wave.

THE CXAIS VIEW

Conclusion

Fix what happens before go-live

The research in this report comes from different firms, different samples and different years, and it points in one direction. Adoption is no longer the problem. **88%** of organisations use AI somewhere.¹⁰ Service leaders are under pressure to do more of it,¹ and they are finding the budget.² What is missing is the foundation that turns spend into return.

That foundation has six parts, and each one has a failure mode that shows up in the data: projects abandoned for want of ready data, business cases built on headcount that never falls, automation layered on broken processes, portfolios chosen by demo, channels designed for a customer who has moved on, and returns nobody agreed how to measure.

None of this requires a new platform. It requires rigour before the build: evidence per workflow, a verdict before a commitment and a number agreed before a pound is spent.

Your 90-day roadmap

Days 1 to 30

Baseline

- Score your programme on the readiness scorecard (page 11).
- Map the data sources each live and planned workflow reads.
- Agree with Finance how AI value will be measured.

Days 31 to 60

Diagnose

- Map in-scope workflows as they actually run, handoffs included.
- Give every workflow a GREEN, AMBER or RED verdict (page 12).
- Name a business owner for each workflow outside IT.

Days 61 to 90

Commit

- Rank the portfolio on value, readiness and risk.
- Write Wave 1 build briefs for GREEN workflows only.
- Move board reporting into P&L terms (page 13).

Methodology and references

This report is a structured synthesis of published research, not a primary survey. We selected sources published between July 2024 and August 2026 from analyst firms, academic institutions and established research bodies, and prioritised studies with disclosed sample sizes. Each finding was mapped to one of the six dimensions in the CXaiS readiness framework. Predictions are labelled as predictions, and where a figure comes from a secondary report of the research, the reference says so. All online sources were accessed in September 2026. Where sources measure slightly different things, we report each on its own terms rather than combining them. The interpretation, verdict criteria and practitioner tools are the view of CXaiS, drawn from delivering enterprise CX technology programmes.

- 1 Gartner (18 Feb 2026). Gartner survey finds 91% of customer service leaders under pressure to implement AI in 2026. Survey of 321 service and support leaders, Oct 2025.
<https://www.gartner.com/en/newsroom/press-releases/2026-02-18-gartner-survey-finds-ninety-one-percent-of-customer-service-leaders-under-pressure-to-implement-ai-in-2026>
- 2 Gartner (26 Aug 2026). Gartner survey finds AI spending by customer service leaders has surged by 38%, despite overall service and support function budgets rising by just 2%. Survey of 199 leaders, Apr to May 2026.
<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>
- 3 Gartner (4 Aug 2026). Gartner survey finds 87% of customers say companies using GenAI for customer service must provide access to a human agent. Survey of 3,566 B2B and B2C customers, Feb to Mar 2026.
<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-agent0>
- 4 Gartner (8 Jul 2026). Gartner survey finds customers are three times more likely to use third-party GenAI than company-provided chatbots for customer service.
<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>
- 5 Gartner (26 Feb 2025). Lack of AI-ready data puts AI projects at risk. Survey of 1,203 data management leaders, Jul 2024.
<https://www.gartner.com/en/newsroom/press-releases/2025-02-26-lack-of-ai-ready-data-puts-ai-projects-at-risk>
- 6 Gartner (25 Jun 2025). Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027.
<https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027>
- 7 Gartner (10 Jun 2025). Gartner predicts 50% of organizations will abandon plans to reduce customer service workforce due to AI. Poll of 163 service leaders, Mar 2025.
<https://www.gartner.com/en/newsroom/press-releases/2025-06-10-gartner-predicts-50-percent-of-organizations-will-abandon-plans-to-reduce-customer-service-workforce-due-to-ai>
- 8 Gartner (2 Feb 2026). Gartner predicts half of companies that cut customer service staff due to AI will rehire by 2027. Survey of 321 service leaders, Oct 2025.
<https://www.gartner.com/en/newsroom/press-releases/2026-02-03-gartner-predicts-half-of-companies-that-cut-customer-service-staff-due-to-ai-will-rehire-by-2027>
- 9 S&P Global Market Intelligence (2025). Voice of the Enterprise: AI and Machine Learning, Use Cases 2025. Survey of 1,000+ respondents in North America and Europe. Figures as reported by CIO Dive, since the full report is subscriber only.
<https://www.ciodive.com/news/AI-project-fail-data-SPGlobal/742590/>
- 10 McKinsey & Company (Nov 2025). The state of AI in 2025: Agents, innovation, and transformation. Global survey.
<https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-state-of-ai-2025>
- 11 Challapally, A., Pease, C., Raskar, R. and Chari, P. (Jul 2025). The GenAI Divide: State of AI in Business 2025. MIT Project NANDA, preliminary findings. Based on a review of 300+ public AI initiatives, interviews with 52 organisations and 153 survey responses from senior leaders. MIT NANDA has no official landing page, so this link is a full PDF copy hosted by AI News.
https://www.artificialintelligence-news.com/wp-content/uploads/2025/08/ai_report_2025.pdf
- 12 Ryseff, J., De Bruhl, B. and Newberry, S. (Aug 2024). The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed. RAND Corporation, RR-A2680-1.
https://www.rand.org/pubs/research_reports/RRA2680-1.html



CXaiS

Take the free AI readiness assessment

Twelve questions, self-service, with your results on screen the moment you finish. Know where you stand before your next pound of AI spend is committed.

- Your maturity level, from Explore to Leading
- A Ready, Constrained or Blocking verdict on each dimension
- An indicative value at stake for your contact volume

cxais.ai/free-assessment



Independent AI in CX consultancy. No platform to sell, no preferred vendor.

CXaiS is a trading name of CXaaS Consulting Ltd, registered in England and Wales, company number 15149940. Registered office: 167-169 Great Portland Street, 5th Floor, London W1W 5PF.