CXaiS | Institute

Framework · November 2026

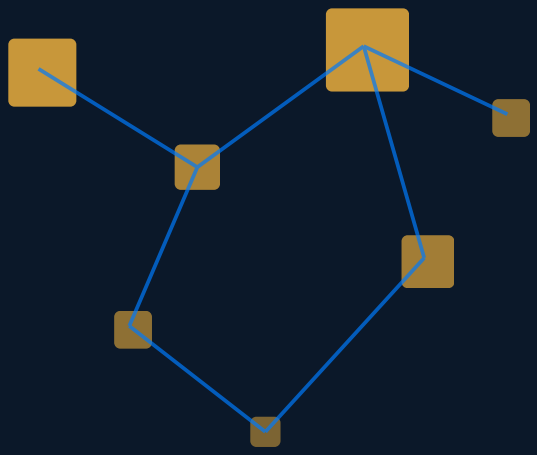

A practitioner framework

The AI fluency framework for CX professionals

Five levels, from first use to setting the standard. What each looks like in a contact centre, how to assess it, and how to tie it to the number.

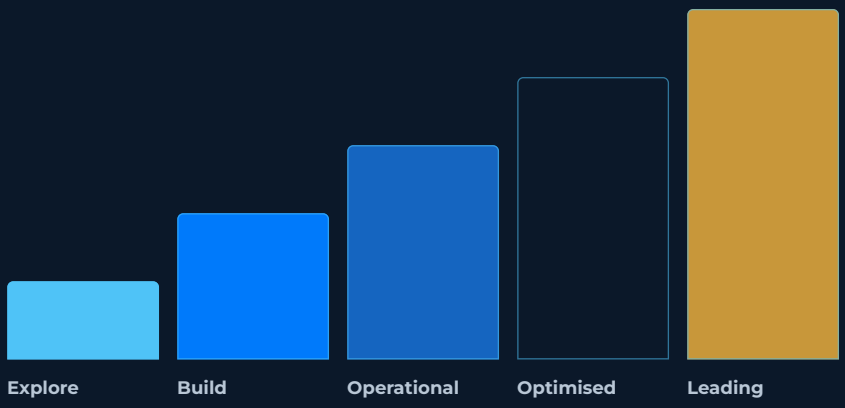

Rachel Lane
Founder and CEO, CXaiS

cxais.ai

Contents

01 Why fluency, not adoption 0
02 The five levels at a glance 0
03 Judgement and safeguards at every level 0
04 Level profiles 0
05 The Operational shift 0
06 Fluency by CX role 0
07 Fluency across the six dimensions 0
08 Assessing fluency 0
09 Measuring fluency 0
10 Rituals that sustain fluency 0
11 Where to start 0
Limitations and sources 0

Organisational readiness and individual fluency, scored in the same five levels.

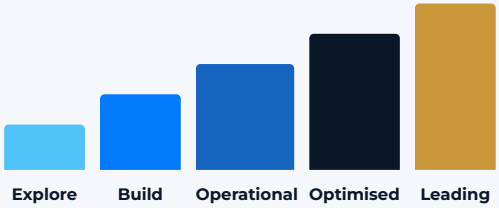

01

Why fluency, not adoption

> “Our agents are using AI” tells a board nothing.

It can mean one team leader summarising calls in a chat window, or a contact centre where every workflow has an owner, a verdict and a number attached to it.

95%

of enterprise generative AI pilots delivered no measurable P&L impact

MIT Project NANDA, 2025

2x

more likely to redesign workflows before selecting AI tools, among companies attributing 5% or more of EBIT to AI

McKinsey, July 2026

The gap matters because access is not the constraint any more. The difference sits in how people, processes and data are prepared, not in licences.

This framework gives CX teams a shared language for that difference. It uses the same five levels as the CXaiS readiness model, Explore, Build, Operational, Optimised and Leading, applied to the individual rather than the organisation. Organisational readiness and individual fluency are scored separately, in the same vocabulary. A CX leader can then answer two questions in the same terms: how ready is the organisation, and how fluent are the people who have to run it?

It is built from practitioner experience inside enterprise CX programs, not a survey. Where a claim is an observation from the field, it is labelled as one.

02

The five levels at a glance

> Each level describes what a CX professional does with AI, not what they know about it. The jump that matters most is Build to Operational.

Exhibit 1 **Five levels of AI fluency, and where the hardest step sits**



Level	What AI is to them	Typical signal
Explore	A curiosity	Tried a tool, went back to the old way
Build	A personal assistant	Saves themselves time most days, within approved tools
Operational	Part of how the work gets done	Checks output, handles exceptions, escalates at the right moment
Optimised	A system to tune	Can show the number moving and say why
Leading	A capability to build in the team	Others reuse their methods, or rely on their judgement

Most CX teams sit across all five at once. In my experience the usual shape is a large Explore and Build population, a thin Operational layer, and one or two people at Optimised or Leading who are carrying far more than anyone realises.

> **Fluency is not authority**
>
> An agent can show Operational fluency inside a workflow someone else designed. Workflow owners carry extra responsibility for design, controls and outcomes, but owning more workflows does not make someone more fluent.



Good fluency includes choosing not to use AI.

Frequency of use cannot tell judgement from enthusiasm.



03

Judgement and safeguards at every level

> Good fluency includes choosing not to use AI.

Frequency of use cannot tell judgement from enthusiasm, so five judgement skills sit underneath every level.

1. **Fit**
   Knowing when AI is the right tool, and when a rule, a search or a human decision is better.
2. **Instruction**
   Giving clear context, constraints and a definition of a good answer.
3. **Scepticism**
   Spotting unsupported answers, missing context and confidence the output has not earned.
4. **Verification**
   Checking output against the authoritative source, not against memory.
5. **Recommend versus act**
   Knowing the difference between AI that drafts an answer and AI authorised to issue a refund or change an account.

Safeguards start at Explore, not at Leading

They grow with the level.

Level	Safeguard expected
Explore	Uses only approved tools, keeps customer data out of anything else
Build	Knows which data is permitted where, and checks before reusing output with a customer
Operational	Verifies, escalates and applies human approval where the workflow requires it

Level	Safeguard expected
Optimised	Tests failure cases systematically and monitors live performance
Leading	Designs the controls, decision rights and stop rules others work within

Five CX failures worth rehearsing

- A fabricated policy answer
- An incorrect refund
- A vulnerable customer the AI did not recognise
- Context lost at transfer
- Customer data pasted into an unapproved tool

Each one is cheap to rehearse and expensive to learn live.

04

Level profiles

Each profile gives the behaviours to look for, what to measure, and the one learning priority that moves someone up.

Level 01 Explore

They have heard the claims and probably tried a tool once. They don't trust the output, and they often have good reason: it was tried on a task it wasn't suited to.

Behaviours	Occasional use, outside core work, often on personal accounts. Scepticism framed as risk or quality concerns.
Measure	Whether they have used an approved tool in the last 30 days.
Learning priority	One task from their own week, done with AI in front of them. Not a generic training module.

Level 02 Build

AI has become a personal habit. They draft responses, summarise contacts, rewrite knowledge articles and prep for calibration sessions faster than before.

Behaviours	Daily use, personal prompts kept in a notes file, good judgement about when output needs checking.
Measure	Frequency and breadth of use across their own tasks.
Learning priority	See their work as a workflow, with inputs, decisions and handoffs, not as a list of tasks.

Level 03 Operational

They have moved AI from their desk into the work. They reliably run their part of an AI-enabled workflow, explain its limits, recognise exceptions and bring evidence for improving it, whether or not they designed it.

Behaviours	Follows the documented workflow, verifies before acting, escalates the cases AI should not handle, logs where it went wrong. Owners also map the workflow before automating it and know who owns each decision.
Measure	Accuracy on a representative task, quality of escalations, and the improvement evidence they contribute. For owners, whether each workflow has a defined outcome measure.
Learning priority	Measurement. Tying a workflow to resolution, cost to serve or retention, not to activity.

Level 04 Optimised

They treat AI workflows as systems to tune. They know which workflows are performing, which are not, and why.

Behaviours	Reviews failure cases, separates containment from resolution, retires workflows that are not earning their place.
Measure	Outcome movement on the workflows they own, and the quality of their explanation for it.
Learning priority	Commercial framing. Expressing the result in the terms a CFO and a board will accept.

Level 05 Leading

They set the standard. Some do it by building capability in others, some as the expert practitioner whose methods and judgement the team relies on. Either way, their workflows, templates and verdicts get reused.

Behaviours	Runs show and tell sessions, writes the playbook, challenges a vendor claim with evidence, coaches Build-level colleagues across the shift.
Measure	Reuse of what they created, and either the colleagues they have moved up a level or the decisions that rest on their judgement.
Learning priority	Governance and sequencing. Deciding which workflows go next, and which wait.

05

The Operational shift

> The move from Build to Operational is where individual productivity becomes organisational capability. It is also where most CX programs stall.

At Build, the gains are real but private. One agent drafts faster, one analyst summarises quicker, and little of it reaches the P&L as measurable benefit. At Operational, AI sits inside the work on purpose. People run their part of the workflow reliably, and the owner knows its inputs, decision points, handoffs and outcome measure.

Three things a training course cannot give on its own

1

Workflow Cartography

The person maps how the work actually happens, including the workarounds, before AI touches it. You cannot redesign what you have not mapped, yet roughly 79% of organisations skip breaking workflows into tasks before deciding what moves to AI (McKinsey, July 2026).

2

The Five Tests of AI Data Readiness

They check whether the data the workflow needs passes reachability, recency, identity, authority and permission at the point of decision. A workflow is only as data-ready as its weakest required test.

3

A Value at Stake line

They can say what the workflow is worth if it performs, in money, before anyone asks.

How to spot a candidate

In my experience the people ready for the shift are rarely the heaviest AI users. They are the ones who already complain about broken handoffs, keep their own process notes, and ask where a number came from.

What leaders have to provide

Time to map a workflow properly, a named owner for each decision in it, and permission to call a workflow RED rather than force it live.



Look for process curiosity, not prompt volume.

The people ready for the Operational shift are rarely the heaviest AI users.



06

Fluency by CX role

> The levels are the same for everyone, but what Operational looks like changes by role. Here is a realistic target for most people in each role.

Role	Realistic target	Operational looks like	Leading looks like
Frontline agent or advisor	**Build to Operational**	Works an AI assist workflow and flags where it gets the answer wrong	Trains peers and feeds failure cases back into the knowledge base
Team leader	**Operational**	Uses AI in coaching and QA on a defined workflow with an outcome measure	Redesigns team routines around AI and spots shift candidates
QA, WFM and planning	**Operational to Optimised**	Runs AI-scored QA or forecasting with a documented calibration check	Separates containment from resolution in every report
CX insight and VoC analyst	**Optimised**	Builds analysis workflows others can rerun	Rebuilds measurement for a world where AI handles most contacts
Knowledge and content owner	**Operational**	Owns the content an AI agent answers from, with review cycles	Sets authority rules for which source wins when two disagree
CX leader, CCO, COO	**Optimised to Leading**	Can defend each AI workflow's number to a CFO	Sequences the portfolio by readiness and value, and calls RED when needed

> Leaders often aim for agents to reach Leading. That is the wrong target.
>
> **Observation from the field**

A contact centre needs most agents at confident Build or Operational, and a small, protected group at Optimised and Leading.

07

Fluency across the six dimensions

> Organisational readiness is assessed across six dimensions. Each one needs a matching skill in people, or the readiness score is a number nobody can act on.

Dimension	The fluency it needs	Question a fluent person can answer
Data quality	Judging whether data is fit for a decision, not just clean	*Does this workflow pass all five data tests at the point of decision?*
Organisational change	Adopting AI in their own work and helping others do the same	*Who in my team is ready for the shift, and what is stopping them?*
Process maturity	Mapping and documenting work before automating it	*Could someone else run this workflow tomorrow from what I have written down?*
Use-case prioritisation	Choosing where AI should go first, and where it should wait	*Which workflow has the highest value and the fewest failing tests?*
Channel architecture	Understanding how AI, assisted and human handling fit together	*Where does this contact hand off, and what context travels with it?*
Commercial alignment	Expressing AI outcomes in money	*What is this workflow worth if it performs, and how will we know?*

Reading the two views together

High readiness, low fluency

The foundation is there but nobody can run on it.

High readiness, high fluency

Ready to scale, workflow by workflow.

Low readiness, low fluency

Start with one workflow and the people who run it.

Low readiness, high fluency

Capable people working around a broken foundation, a retention risk you will feel first in your best people.

08

Assessing fluency

> Assess what people can demonstrate on a real task, within their role and their opportunity to practise. A manager's estimate tends to reward confidence and visibility, and access to better tools is not a higher level.

Evidence	What it demonstrates	Level it confirms
Completes a representative task from their own week	Practical application	**Build**
Spots a deliberately flawed AI output	Verification and scepticism	**Build to Operational**
Explains when they would escalate, and to whom	Awareness of limits	**Operational**
Shows a repeatable method someone else could follow	Operational reliability	**Operational**
Provides before and after evidence on a workflow	Optimisation	**Optimised**
Points to methods or judgement others now rely on	Standard setting	**Leading**

Three inputs, one rule

Short self-assessment + **Work sample on a real task** + **Calibrated reviewer**

Combine the three, with the reviewer calibrated against the level profiles. Where they disagree, the work sample wins.

09

Measuring fluency

> Measure fluency at the workflow, not the login. Usage data shows who has started, not who has made the Operational shift.

Exhibit 2 **Three layers of measurement, each answering a different question**

Usage Explore to Build	Active use of approved tools, breadth across tasks.	*Is AI part of their day?*
Ownership Build to Operational	Workflows run and documented with AI, each with a named outcome measure.	*Is AI part of the work?*
Outcome Operational to Leading	Movement in resolution, repeat contact, customer effort, complaints, successful handoff to a human, accessibility, cost to serve, retention or revenue, plus reuse by others.	*Is AI earning its place?*

The outcome layer is where fluency meets Value at Stake. Every Operational workflow should carry a Value at Stake line from the start, so the fluency program reports in the same currency as the AI investment it supports.

Be precise about what counts as value

Time saved → **Usable capacity** → **Realised financial benefit**

Time saved becomes usable capacity only if the time is redeployed, and capacity becomes realised financial benefit only if it shows up in cost, revenue or retention. Net value accounts for checking and correction time, implementation, running costs and adoption. Set a baseline and a comparison before go-live, because a number moving after launch does not by itself prove the AI moved it.

A warning on containment

When AI handles most contacts and surveys only go to contacts a human handled, you are sampling the exception path. A fluent CX professional checks the sampling design before reporting a rising CSAT as proof the AI is working.



Measure fluency at the workflow, not the login.

Usage tells you who has started. Outcomes tell you who has arrived.



10

Rituals that sustain fluency

> Fluency decays when the tools change and the habits don't. Five rituals keep it moving, each short enough to survive a busy quarter.

Weekly

Failure review, 30 minutes

The team looks at five contacts where AI got it wrong. The aim is to fix the workflow, not to blame the model.

Monthly

Workflow show and tell

One Operational or Optimised owner walks through a workflow, its map and its number. Build-level colleagues see what the shift looks like in practice.

Quarterly

GREEN/RED review

Every live AI workflow gets a verdict on whether it is performing. A workflow missing its commercial target goes RED to be sequenced and fixed, not abandoned. A workflow causing customer harm or taking actions it is not authorised to take is paused immediately, by a named person, and restarts only when the fix is tested. The quarterly review sits on top of live monitoring, never instead of it.

Always on

A shared workflow library

Maps, prompts, templates and calibration notes in one place, owned by a Leading-level person. Reuse is the clearest Leading signal.

Every release

Vendor release check

When a platform ships a new AI feature, someone asks which existing workflow it changes before anyone switches it on.

11

Where to start

> Start with one workflow and the people who run it, not a company-wide training plan.

1 **Place your team**

Use the assessment method, one real task per person, to count how many people sit at each level. Start with the team that runs your first workflow, not the whole function.

2 **Find two shift candidates**

People at Build with process curiosity, using the signals in the Operational shift section.

3 **Pick one workflow**

High volume, a clear outcome measure, and data that passes every test the workflow requires. Where a required test fails, fix it before go-live.

4 **Map it, test it, value it**

Workflow Cartography, the Five Tests and a Value at Stake line, done with the candidates rather than for them.

5 **Run the rituals for one quarter**

Then recount the levels and report the workflow's number alongside the fluency shift.

See where your organisation sits on the same five levels

The free AI Quick Assessment scores readiness across the six dimensions, so you can set organisational readiness and team fluency side by side.

cxais.ai

Limitations and sources

This is a practitioner framework built from experience inside enterprise CX programs. It has not yet been validated as an assessment instrument, and the CXaiS Institute will test it against real team data before publishing benchmarks.

The MIT Project NANDA figure is a widely cited headline from a 2025 study of enterprise generative AI pilots. Its sample and definition of measurable impact have been debated, so it is best read as a direction of travel rather than a precise rate.

Sources

1. MIT Project NANDA, *The GenAI Divide: State of AI in Business 2025*.
2. McKinsey & Company, The operating model advantage: Why AI winners are rewiring their organizations, 7 July 2026.

About CXaiS

CXaiS is an independent AI in CX consultancy. We build the commercial foundation, ROI framework and implementation plan that makes AI investment in customer experience perform. We have no platform to sell. The CXaiS Institute publishes the research and frameworks behind that work.

CXaiS

Where does your organisation sit?

Take the free AI Quick Assessment. Twelve questions across the six dimensions, and an instant report on where you stand.

cxais.ai

Talk to us: enquiries@cxais.ai

CXaiS Institute is the research imprint of CXaiS. CXaaS Consulting Ltd, registered in England and Wales, company number 15149940. Registered office 167-169 Great Portland Street, 5th Floor, London W1W 5PF. © 2026 CXaaS Consulting Ltd.