Typewise

Beyond the Bot: A CX Maturity Playbook for Multi-Agent AI

AI Adoption in Customer Experience (CX) Whitepaper (February 2026)

https://typewise.app

Executive Summary

Enterprise AI in customer experience is moving out of the lab and into production. Across our 36 C-level interviews, one pattern is clear: the majority of value realized today comes from assistive AI, systems that draft, guide, and support human agents, while semi-autonomous and guardrailed autonomous deployments are emerging in targeted, carefully controlled workflows.

Autonomy does not arrive everywhere at once. It appears first where three preconditions align:

1. **API coverage** for the end-to-end action,
2. a credible **evaluation harness** that defines "good enough" behavior, and
3. an **operating model** that makes risk observable, reversible, and owned.

From this research, five conclusions stand out:

1. **Adoption is uneven, but real.** Industries with structured products and cleaner integration surfaces show a higher share of organizations in scaled production, while integration-heavy or highly regulated sectors cluster at pilot or limited production. Larger enterprises advance faster when a central platform team owns policy, evaluation, and observability.
2. **Autonomy lands where actions are verifiable and reversible.** Today, autonomy concentrates in chat, email, and in-app channels for bounded tasks such as cancellations, refunds, and profile changes. Full voice automation is rarer and appears where identity, policy, and failure modes are highly engineered.
3. **Make-or-buy is a per-workflow decision.** Scale and sovereignty push organizations toward make or hybrid approaches; speed and non-differentiating tasks favor buy or partner-led pilots. The most resilient architectures decouple the agent layer from channels and CRM platforms, preserving policy control and model choice.
4. **Maturity, not model choice, determines what is possible.** Governance and evaluation maturity are the strongest predictors of safe autonomy; stack orchestration and data readiness determine where autonomy can be applied.

5. **A predictable playbook works.** High-volume, API-ready workflows in Support move first; Professional Services follows with structured tasks; Sales benefits from assistive coaching before any autonomous actions. Progress is governed by explicit promotion gates up the Agentic Ladder, rather than by enthusiasm or vendor promises.

Our recommendation: Treat autonomy as a license the system must qualify for. Start by using assistive AI in one high-volume, API-ready workflow (for example, subscription cancellations or simple plan changes). Use this phase to prove safety and value, and to build your evaluation and policy controls as platform capabilities. Only then grant the AI more action rights: moving to semi-autonomous and, in time, guardrailed autonomous operation when your maturity gates are clearly met.

Research Methodology:

- Sample: 36 structured interviews with C-level leaders across Europe.
- Industries: software, IT services, internet, retail & e-commerce, gambling & casinos, travel & hospitality, automotive, facility management & real estate, logistics & transportation, insurance, public sector, utilities.
- Roles: Chief Customer Officers, Chief Digital/Technology Officers, Heads of CX, and equivalent.
- Approach:
    - Transcripts coded against a normalized schema (enums; 0–5 maturity rubrics).
    - Records joined with firmographics (industry, employee band, country/region).
    - Aggregations built for stage, posture, sourcing, channels, use cases, and maturity.
    - Quotes scored for signal, originality, attribution strength, and thematic fit.
    - A full interviewee list available in the Appendix.

1. Introduction and Context

Customer experience has become the primary interface between an enterprise and its customers. Service expectations are high, tolerance for friction is low, and cost pressures are persistent. Traditional automation such as IVR trees, scripted bots, rigid self-service flows, has delivered incremental gains, but has largely exhausted its promise.

Generative and agentic AI raise the ceiling. They can synthesize knowledge, reason across systems, and orchestrate actions. But they also introduce qualitative risks: hallucinations, inconsistent tone, policy violations, and opaque decision-making. For C-level leaders, the central question is no longer “Can we use AI in CX?” but “How do we operationalize it safely and at scale?”

The organizations we interviewed are converging on similar answers:

- They prioritize **high-volume, API-backed, reversible** workflows.
- They invest early in an **evaluation harness** and **policy enforcement**, not only in models.
- They introduce autonomy gradually, via an explicit **Agentic Ladder**.
- They treat AI not as a feature but as a **platform capability** that underpins multiple use cases.

The sections that follow describe the adoption landscape, the maturity operating system, and the frameworks we derive from the interviews.

> *“AI helps us streamline the business and gain efficiency, but in the end, everything comes back to the client experience. We only call something a success if the end client feels the improvement.”*
>
> Ross Abbate, CEO, Macro

2. Adoption Patterns: What Is Live Today

2.1. Where organizations are on the adoption curve

Across the panel, most organizations have moved beyond experimentation. They sit somewhere along a continuum from no live deployments to scaled production. The distribution is not uniform.

- In some industries, particularly those with digital products and modern integration surfaces, leaders report a meaningful share of use cases in scaled production.
- In others, especially those with complex legacy stacks or tight regulatory oversight, AI remains concentrated in pilots or limited production deployments.

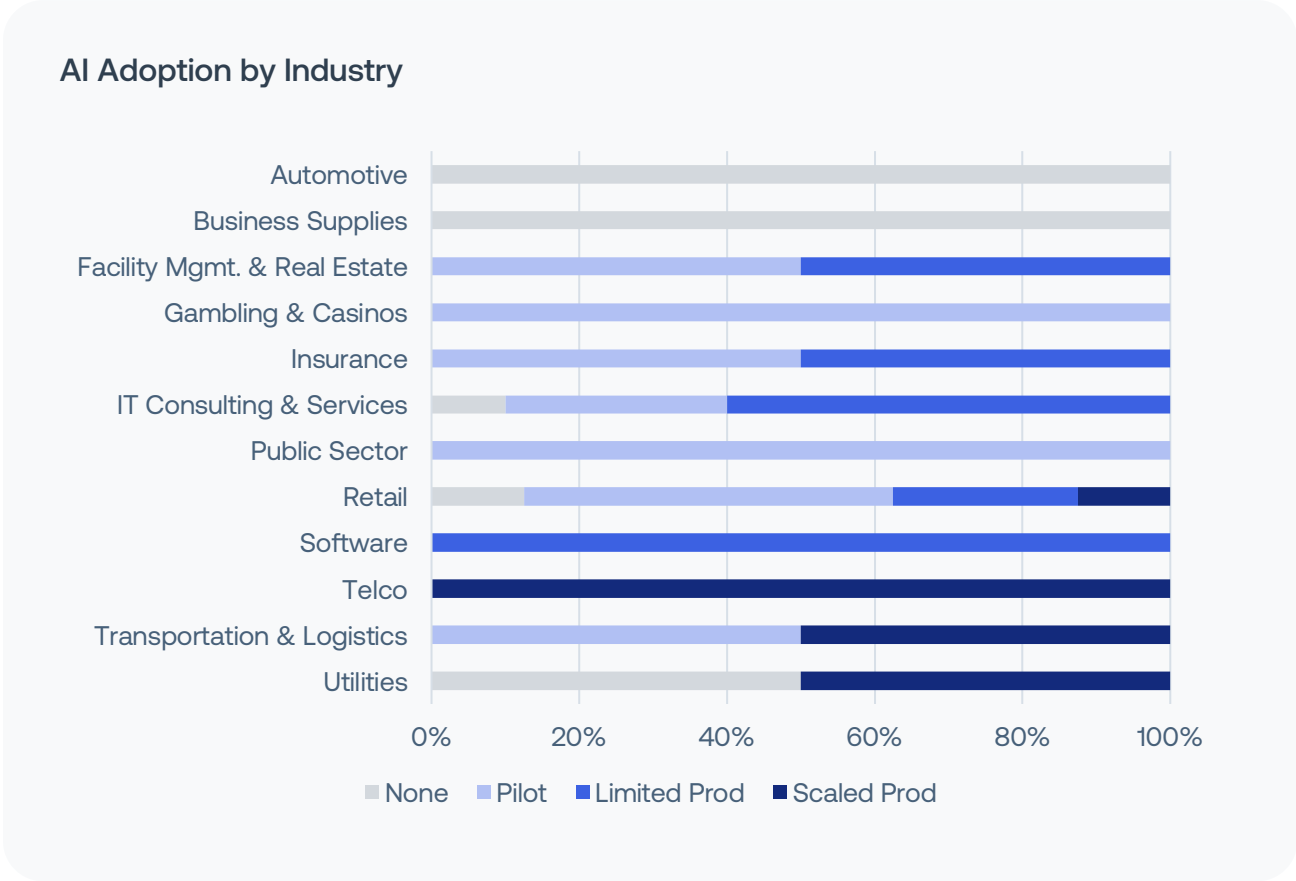

2.2. Agentic posture by channel

The **agentic posture**, what the AI is allowed to do, varies significantly by channel:

- In **chat and email**, assistive AI is pervasive. Agents receive suggested replies, knowledge retrieval, and next-best actions. Semi-autonomous and guardrailed autonomous deployments appear where API coverage is strong and actions are reversible.
- In **in-app** contexts, some enterprises already run guardrailed autonomy for clearly bounded tasks, supported by UI confirmations and policy checks.
- In **voice**, autonomy is more constrained. Leaders cited latency, intent disambiguation, and failure-mode design as key concerns.

Agentic Posture by Channel

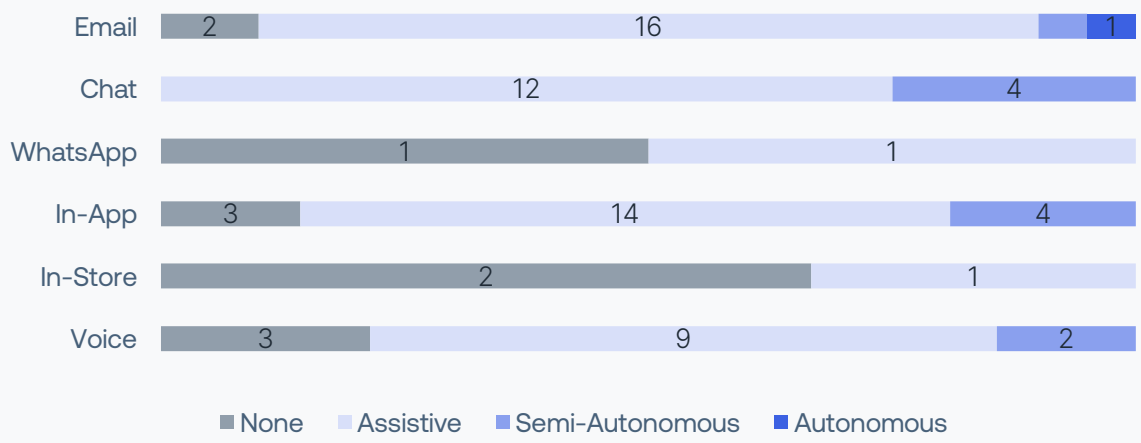

> *"The real opportunity is in email and chat, where AI can handle the bulk of simple requests and leave the complex cases to our agents."*
>
> Michal Paschialis, CEO, IAI Group

2.3. Use-case penetration by industry

Across industries, a consistent set of use cases dominates:

- **Agent assist**: knowledge retrieval, summarization, drafting, classification.
- **Self-service**: FAQ-style interactions and transactional flows where policy logic is clear.
- **QA and coaching**: evaluating agent responses against policy and tone and suggesting improvements.
- **Voice bots**: more prevalent where call intents and IVR stacks are well understood but still lagging relative to chat/email.

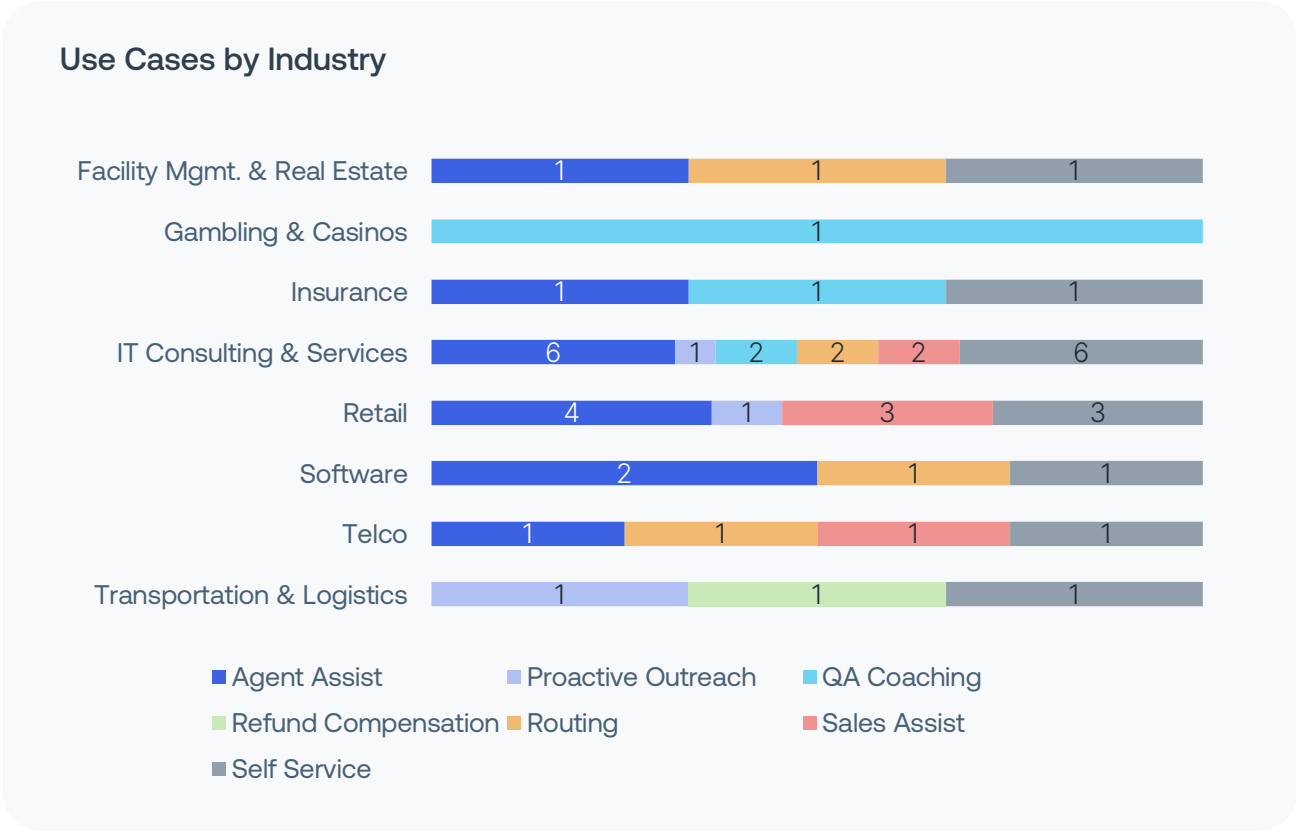

Industries with codified policies and repeatable transactions, such as subscription businesses or certain travel contexts, move faster on self-service. Others use AI primarily in assistive and QA roles while they improve their knowledge base and integration layer.

2.4. Blockers and enablers of progress

Interviewees described a consistent set of **blockers**:

- **Integration debt**: key workflows lack robust APIs or have fragmented identity and entitlement models.
- **Data fragmentation**: knowledge is scattered across tools, with no single source of truth.
- **Brand and trust risk**: concern about tone, hallucinations, and inconsistent experiences.
- **Legal and regulatory constraints**: especially in financial services, healthcare, and public sector environments.
- **Internal trust and change fatigue**: reluctance from agents and middle management to accept AI-driven changes.

They also identified clear **enablers**:

- **Central platform teams** responsible for orchestration, policy, and evaluation, rather than siloed projects.
- **Clean API surfaces** for core workflows, with explicit action rights and auditability.
- **Early investment in evaluation harnesses**, including representative offline tests and online metrics.
- **Executive sponsorship** that links AI initiatives to concrete CX and cost objectives.

> *"If we want to use AI, we have to show clearly how it benefits members and that it doesn't compromise their trust in us."*
>
> Steve Finlan, CEO, The Wine Society

Blockers vs. Enablers

Blockers

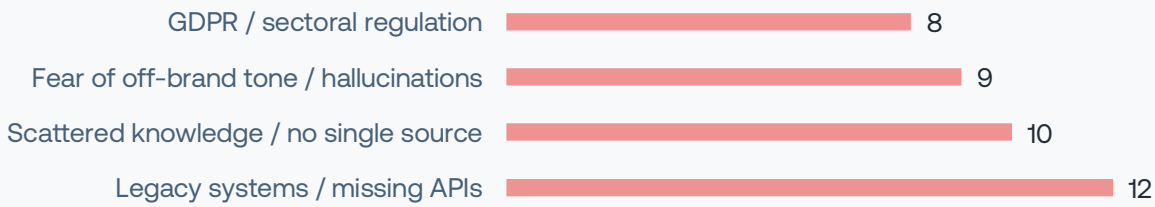

Enablers

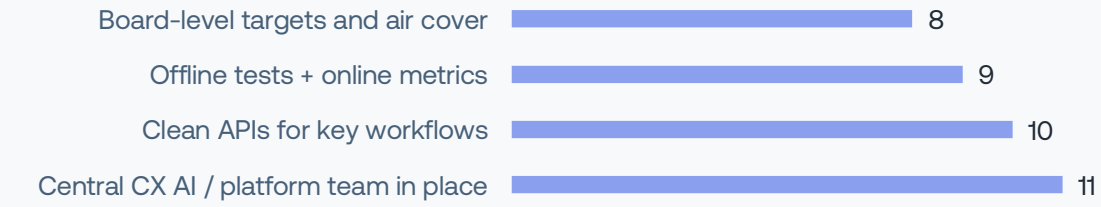

Implication: If your AI program is stalled, the most useful question to ask is *not* "Which model are we using?", but "Which of these blockers have we actually removed?"

> *"We have very clear standards for data management and data governance, including which tools we allow into the group at all."*
>
> Karsten Rech, CIO, Vonovia

3. CX AI Maturity OS: Operating System for Progress

The CX AI Maturity OS is a six-dimension framework that describes whether a workflow is ready to move up the Agentic Ladder. It links strategy, data, governance, architecture, evaluation, and operating model into a single view.

3.1. The six dimensions

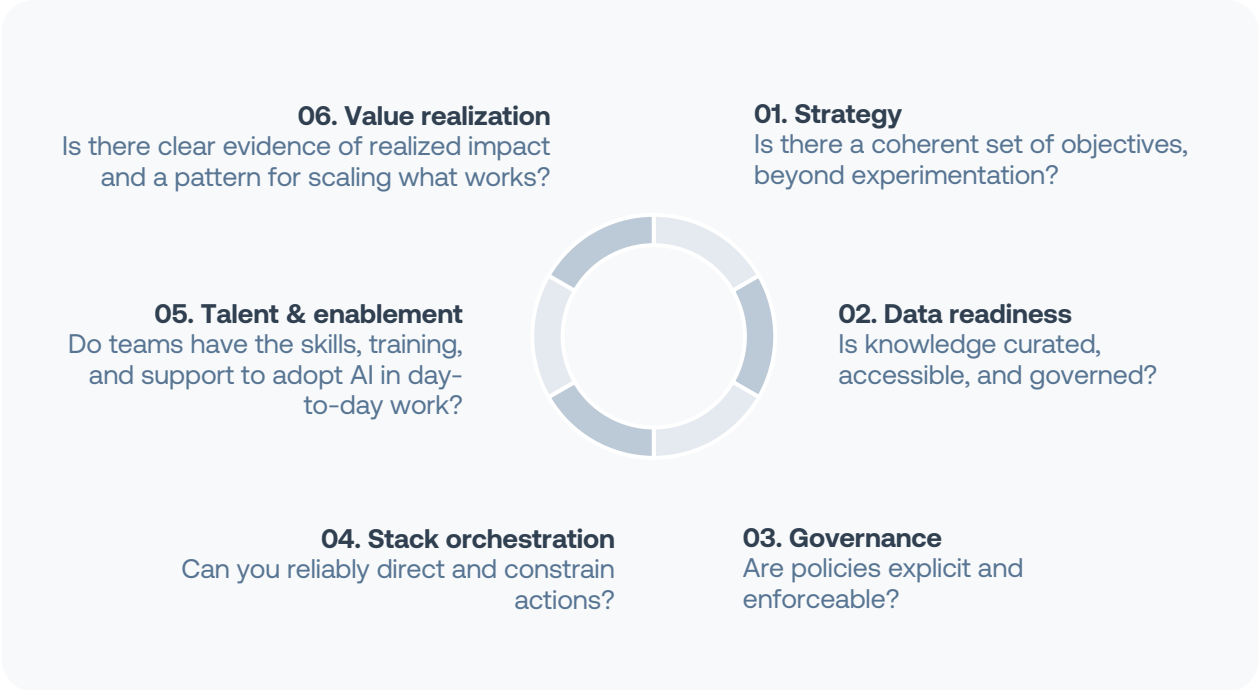

Each dimension is scored on a scale from 0 to 5, with anchor descriptions. The detail appears in Appendix A; the key idea is simple: do not promote a workflow to higher autonomy unless its maturity scores justify the risk.

> *"Back in 2023, we took an architectural approach rather than a single-LLM experiment. Instead of "plugging in ChatGPT," we designed for a multi-agent AI setup from day one (including our own SLM), assuming multiple specialized agents would serve different business domains. This required not only a robust data structure, but also clear information safeguards, access controls, and governance to ensure scalability and futureproofing across departments."*
>
> Emilija Frew, Head of Tech, Pierce Group

3.2. Cross-industry maturity patterns

When we aggregate the six maturity dimensions by industry cluster, a clearer pattern emerges.

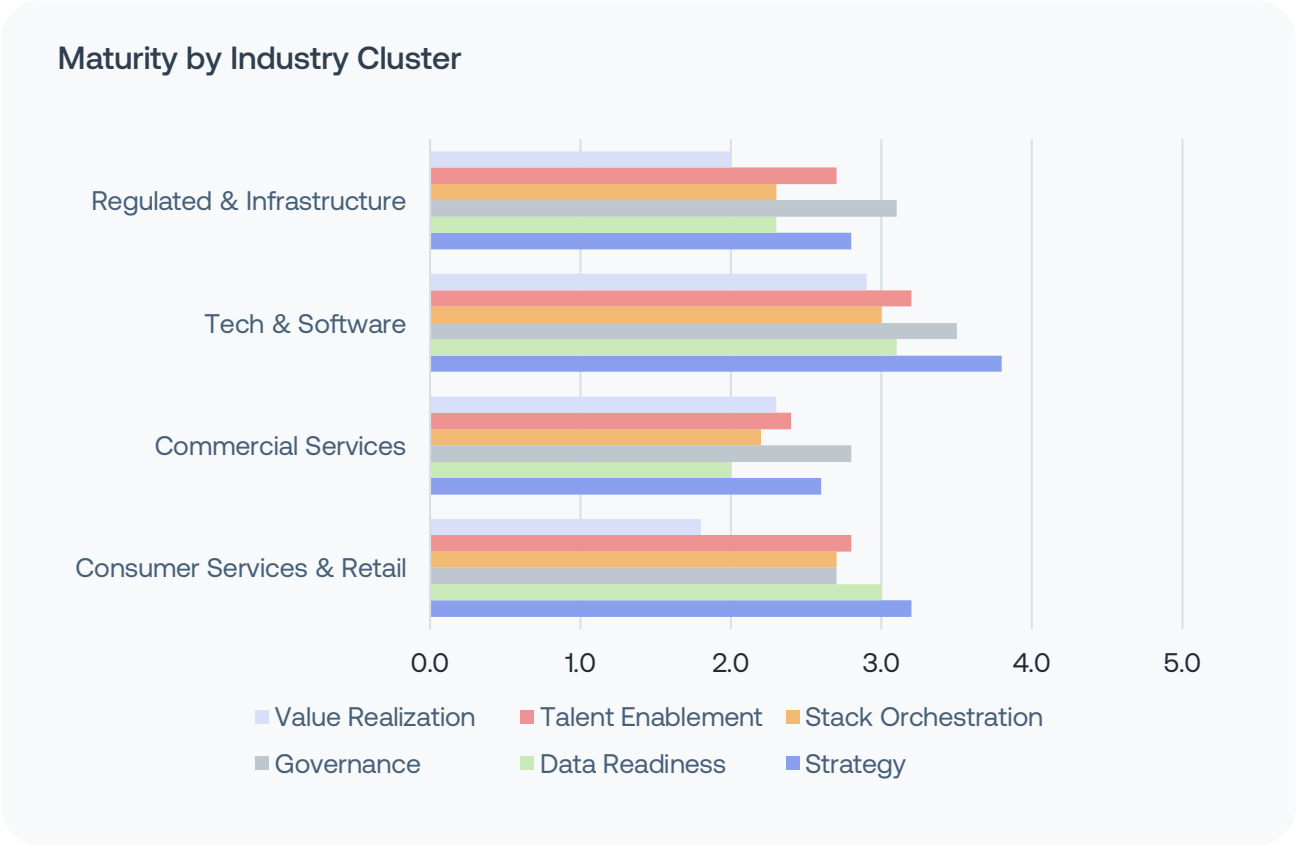

First, no sector is “finished”. Even the most advanced industries sit in the “good, not yet world-class” range on average. That is consistent with a market where CX AI is still in early deployment, and where most organizations are building their operating model as they go. Among the clusters:

- **Tech & Software** sits slightly ahead overall and tend to score higher on strategy, stack orchestration, and talent enablement. They are more likely to have a central platform team, clearer roadmaps, and API-ready systems, which allows them to move faster when they decide to scale a workflow.
- **Regulated & Infrastructure** sectors show stronger governance and risk maturity than many consumer sectors, but more modest scores on value realization. They are investing in controls and evaluation yet are cautious in translating those capabilities into broad automation and measurable results.

- **Consumer Services & Retail** are mixed. Many players have solid strategic intent and strong frontline urgency, but their scores on data readiness and stack orchestration are lower, reflecting fragmented systems and legacy CX stacks. This matches comments about “wanting to move faster” but being slowed by inconsistent data and integration.
- **Commercial Services** often sit in-between. They usually have clearer owned processes and relationships, but still face integration challenges across CRMs, field-service tools, and finance systems. Their maturity scores suggest that once orchestration and data gaps are addressed, they are well placed to scale AI beyond assistive use.

Across all clusters, two themes are consistent:

1. **Governance and evaluation maturity track safe autonomy.** Industries that have invested in policy engines, auditability, and evaluation harnesses are the ones running early guardrailed autonomous flows.
2. **Data and integration are the limiting factors. Data readiness** and **stack orchestration** systematically lag strategy and governance, especially in organizations with multiple legacy systems and inconsistent schemas. Until those gaps close, many leaders are deliberately keeping AI at assistive or narrowly semi-autonomous levels, regardless of how strong the underlying models are.

This pattern supports the central message of the Maturity OS: it is not enthusiasm or model choice that separates industries, but the extent to which they have turned governance, evaluation, and orchestration into repeatable capabilities that can be applied across workflows.

> *“We’re starting with very concrete use cases. For support, AI helps us summarize tickets, suggest responses and highlight risks, but the agent stays in control. Over time we’ll let the system handle more of the routine interactions itself.”*
>
> Jon-Gunnar Aasen, CEO, Miles

3.3. Using maturity as a decision gate

In practice, the CX AI Maturity OS is not a scorecard for the organization as a whole. It is a decision gate for each workflow. Leading teams use it in three simple steps.

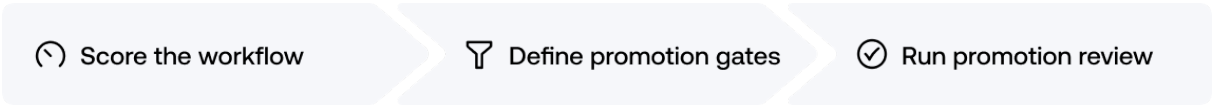

Score the workflow	Define promotion gates	Run promotion review
Assess six dimensions on a 0–5 scale.	Set minimum scores for promotion aligned with risk appetite.	On a regular cadence, decide per workflow: stay, promote up the Agentic Ladder, or roll back.

1. Score each workflow, not the organization

For every candidate workflow, such as cancellations, billing adjustments, or high-stakes complaints, the team scores the six dimensions on a 0–5 scale: strategy, data readiness, governance & risk, stack orchestration, talent & enablement, and value realization. The result is a compact view of where the workflow is strong and where it is fragile.

2. Define promotion gates

Executives then agree explicit “promotion gates” for moving up the Agentic Ladder. For example:

- Governance ≥4/5 (policies, ownership, auditability in place)
- Stack orchestration ≥3.5/5 (reliable APIs, clear action rights, rollback tested)
- Value realization ≥3/5 (baseline metrics and early impact demonstrated)

Only workflows that meet these gates are even considered for promotion from assistive to semi-autonomous, or from semi-autonomous to guardrailed autonomy.

3. Run periodic promotion reviews

On a regular cadence (typically every 4–6 weeks), a cross-functional group reviews each workflow against its maturity scores and recent performance. The decision for each workflow is binary and explicit:

- **Stay** at the current autonomy level while gaps are addressed.
- **Promote** one rung up the Ladder, with updated guardrails and monitoring.
- **Roll back** if performance, risk context, or controls have deteriorated.

This turns AI deployment from an act of belief ("we think the bot is ready") into an act of **controlled progression**. Workflows earn more autonomy based on evidence and agreed thresholds and can lose it if the evidence changes.

> *"For a lot of incoming questions, we expect AI to propose the answer, but with a human validator behind it. That's how we think about the first stage: AI does the work, yet people stay in charge of the decision."*
>
> Soren Meyer, CEO, ZeroNorth

4. Make–Buy Decision: Structuring Investment Choices

The question “Should we make or buy?” is too blunt for the reality of CX. In practice, enterprises end up doing all four: making, buying, running hybrid architectures, and using partners. The real decision is which posture fits each workflow, at this point in time, given your risk, data, and integration constraints.

The **Make–Buy Decision Matrix** organizes that choice along two axes:

1. **Data sensitivity and sovereignty (horizontal):** How constrained are prompts, logs, and outputs by regulation or internal policy? How severe are the consequences if they leak or are mishandled?
2. **Integration debt and time-to-value (vertical):** How many systems must be orchestrated? Are there clean APIs? How quickly do you need the workflow in production to prove value?

Make–Buy Decision Matrix

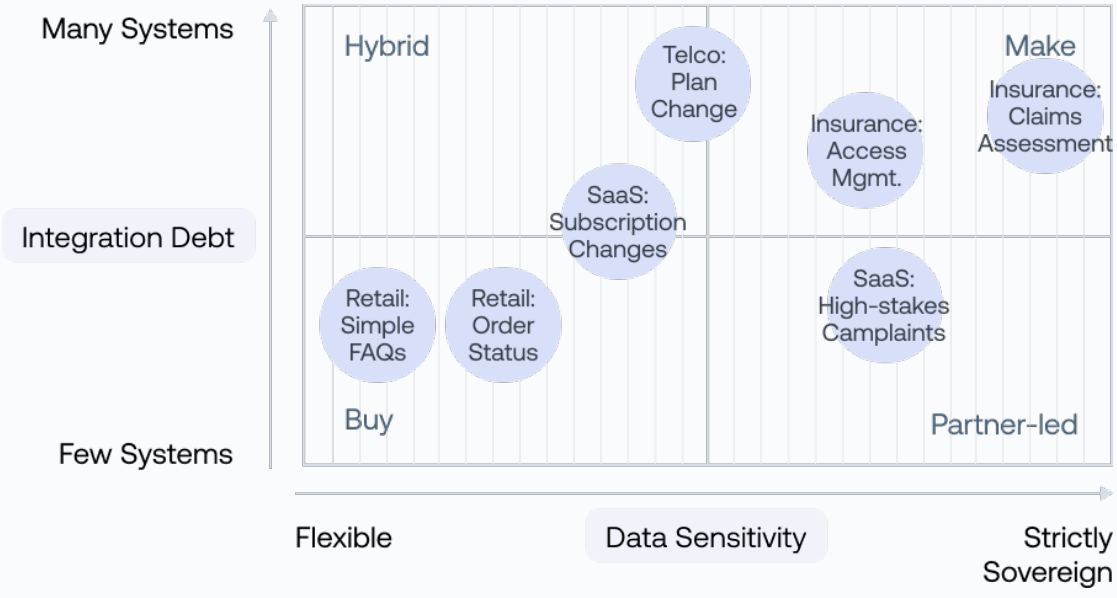

Plotted on these axes, typical CX workflows naturally cluster into **Buy**, **Hybrid**, **Make**, or **Partner-led** zones. The exhibit provides a short description of each quadrant; this section focuses on how leaders actually use the matrix.

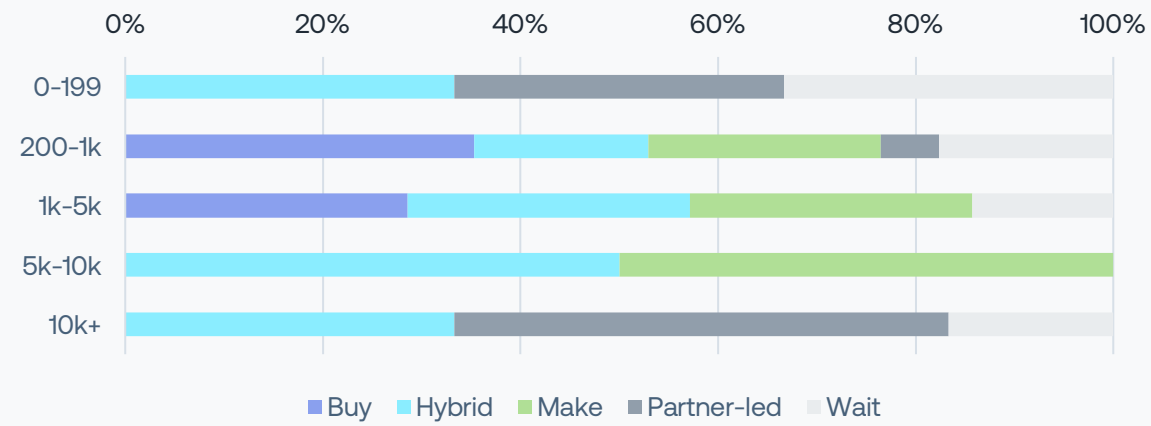

4.1. Sourcing patterns in practice

From our interviews, three pragmatic patterns emerged:

Buy for commodity, low-risk flows

Leaders buy off-the-shelf where:

- Workflows are non-differentiating (e.g., basic FAQs, generic support widgets).
- Data sensitivity is low and existing SaaS controls are sufficient.
- The main objective is speed to value rather than deep customization.

> *"Teams spend far too much time defining requirements and building proofs of concept. If someone has already solved a problem and it comes out-of-the-box from a platform, we test that first."*
>
> Ida La Spisa, CIO, Telia

Hybrid as the default for serious CX

When the workflow touches core customer data, complex entitlements, or critical brand interactions, executives prefer to keep the agent, policy, and evaluation layer in-house, even if they lean on vendors for channels, telephony, or models. This gives them:

- Control over what actions are allowed and when.
- The ability to swap or mix models without redesigning the entire stack.
- A single place to enforce governance, logging, and evaluation.

Make when sovereignty and differentiation align

For a smaller set of workflows, typically those that define the customer relationship or are tightly regulated, leaders opt to build. These are places where:

- Data cannot leave their own environment.
- The experience is fundamentally differentiating.
- They are prepared to invest in a platform team and long-term ownership.

> *"A good AI strategy is usually 'make and buy' – taking the best of both worlds. We build where it gives us a real competitive edge, and we buy where others are already excellent."*
>
> Hüseyin Dogan, COO, IONOS

Partner-led sits across this: executives use partners to accelerate pilots and change management when their own teams are bandwidth-constrained, then deliberately transition successful patterns into their platform operating model.

4.2. Applying the matrix per workflow

In practice, leaders use a simple four-step approach for each workflow:

1. **Score each candidate workflow**
    - Sovereignty / sensitivity: 1 (low) to 5 (very high).

- Integration complexity: 1 (few, clean systems) to 5 (many, fragmented systems).
- Required time-to-value: how soon you need a credible business result.

2. **Place the workflow into a quadrant**: Using the matrix, mark each workflow where its scores suggest it belongs. A cancellation flow in a digital-only product might sit in **Buy**; a refund plus policy exception flow on regulated accounts might sit in **Hybrid** or **Make**.
3. **Check consistency with your maturity**
   - If governance and evaluation are weak, even "Make" candidates might need a **Partner-led** phase first.
   - If stack orchestration is not ready, a "Hybrid" plan might need to start with a **Buy** configuration while you build your platform capabilities.
4. **Revisit annually**: As your maturity increases and vendor offerings improve, some workflows will naturally migrate across quadrants. Buy flows may become hybrid; partner-led flows may come fully in-house.

4.3. Reference architecture: decoupling the agent layer

Sourcing choices are not only commercial; they are architectural. The most resilient organizations in our panel are converging on a similar pattern:

1. **Channels remain diverse.** Customers arrive via chat, email, voice, in-app, and social. Organizations continue to use best-of-breed telephony, messaging, and CRM tools where it makes sense.
2. **Agent and orchestration logic sits in its own layer.** Instead of embedding "the bot" inside a single CRM or contact-center suite, leaders introduce an AI-first orchestration layer that:
   - Interprets requests and maintains conversational state.
   - Applies policies and the Agentic Ladder (what actions are permitted in this workflow).
   - Selects tools and models and calls systems of record through APIs.

3. **Systems of record stay the source of truth.** Customer data, orders, tickets, and payments continue to live in CRMs, ERPs, and line-of-business systems. The agent layer reads and writes via explicit, audited APIs.
4. **Models are swappable.** Large language models and other AI components are treated as pluggable, not hard-wired into any single channel or system.

Reference architecture

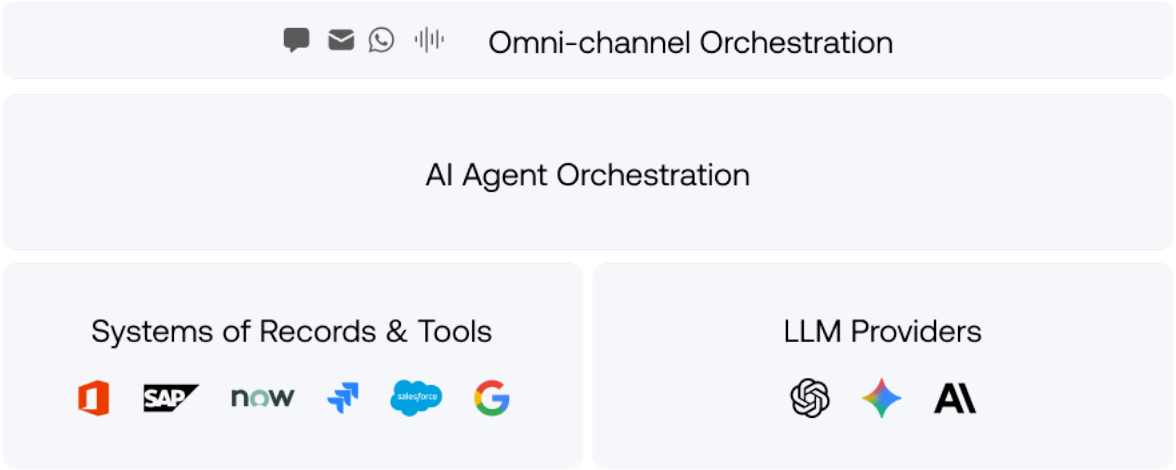

One C-level leader described it as "laying an agent layer *on top* of existing systems, rather than buying yet another monolithic platform that wants to own everything." Traditional vendors are well placed to provide channels and tickets; they are not always the best place to centralize policy, orchestration, and evaluation.

Architecturally, this gives executives three advantages:

- **Policy control in one place.** Promotion gates, action rights, and audit trails are enforced consistently across channels.
- **Vendor flexibility.** You can evolve CRM, contact-center, and model choices without redesigning every workflow.
- **Faster reuse.** Once one workflow is proven from a given channel, the same orchestration can be reused in others with minimal changes.

5. The Agentic Ladder: Making Autonomy Safe

Most senior leaders tell us they are uncomfortable with phrases like “we’ve launched an AI bot.” That wording hides the question they care about most: **“Exactly what is this system allowed to do, and who decided that?”**

The **Agentic Ladder** is a simple but powerful answer to that question. It describes, in plain language, the level of autonomy the AI has today, and what must be true before it is allowed to do more. Instead of a binary “bot / no bot”, you get a controlled progression:

- Start with assistive behavior that cannot harm customers.
- Demonstrate value and reliability under observation.
- Only then, and only where controls are in place, grant the system limited action rights.
- Reserve full guardrailed autonomy for workflows that have clearly earned it.

For boards and risk owners, the Ladder provides transparency. For product and engineering teams, it becomes a **design contract**: “If we want to move this workflow up one rung, here is what we must build and prove.”

Agentic Ladder

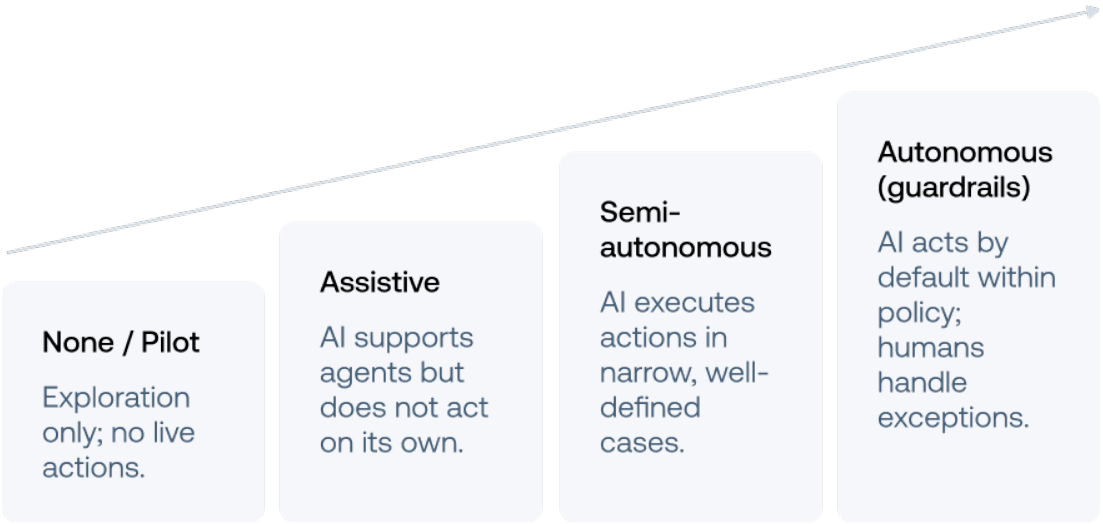

5.1. Why explicit autonomy levels matter

Across the interviews, we heard two recurring failure modes:

- **Over-promising autonomy**: a pilot is described as “fully automated”, but in reality, still relies on humans correcting errors behind the scenes. Trust erodes when stakeholders discover the gap.
- **Silent drift**: a system gradually takes on more actions, new tools, new flows, without a conscious decision or updated risk assessment.

Both problems stem from the same cause: no shared language for “what the AI is allowed to do” and no formal process for changing that scope.

The Agentic Ladder addresses this by:

- Giving each level a name, behavior pattern, and owner.
- Tying advancement to maturity gates (governance, evaluation, orchestration, data).
- Making it explicit that autonomy is **reversible** if performance or controls deteriorate.

> *“Our AI layer pulls data from our systems and prepares an answer, but the customer service agent still has the final approval. That way we get the speed of automation while keeping a human in control of the outcome.”*
>
> Pavlo Antonov, CTO, SelectSpecs

5.2. The four autonomy levels

Level 1: None/Pilot

- The AI is being explored in a safe environment—often on historical data or in shadow mode.
- It does not interact with customers or take actions in production.

Objective: understand feasibility, data quality, and integration gaps.

Level 2: Assistive

- The AI supports agents but does not act on its own.
- It drafts responses, retrieves knowledge, summarizes conversations, and suggests next-best actions.
- Agents remain fully in control; they accept, edit, or ignore suggestions.

Objective: prove that the system can consistently lift productivity and quality without increasing risk.

This is where most organizations in our panel are operating today.

Level 3: Semi-autonomous

- The AI is allowed to execute actions directly within a clearly defined scope (for example, refunding small amounts, applying safe credits, or scheduling simple appointments).
- It does so only when certain conditions are met: high confidence, clear identity, and policy checks passed.
- Outside that scope, it falls back to assistive behavior and human approval.

Objective: move a subset of volume to AI-driven execution while keeping humans in the loop for edge cases and sensitive scenarios.

Level 4: Autonomous (with guardrails)

- For a bounded class of workflows, the AI acts by default.
- Guardrails—policy checks, confidence thresholds, monitoring, and kill-switch—ensure that actions remain within agreed limits.
- Humans are involved through sampling, QA, and exception handling, not by default on every interaction.

Objective: achieve meaningful automation in areas where the risk is acceptable, and the controls are robust.

5.3. Promotion gates for autonomy

The key principle is that no workflow moves up the Ladder by default. It earns the move by meeting defined criteria:

Governance & risk

- Documented policies for what the AI may say and do in that workflow.
- Audit trails for prompts, retrieved data, and actions.
- Clear ownership for risk, escalation paths, and incident management.
- Regular review by a cross-functional group (CX, IT, Legal/Compliance).

Evaluation

- A representative offline test set for that workflow, with metrics and thresholds agreed in advance.
- Online monitoring of quality, safety, and customer outcomes.
- Evidence that performance is stable over time and across volumes.
- Regression tests when models, prompts, or policies change.

Stack orchestration

- APIs or other reliable mechanisms for the AI to perform allowed actions.
- Explicit action rights (who/what may do what, under which conditions).
- Tested rollback procedures for misfires or misroutes.
- Observability: the ability to inspect what happened, step by step.

Data controls

- Classification of the data involved (including PII and sensitive fields).
- Masking/tokenization and retention policies applied consistently.
- Clarity on where data and logs are processed and stored.

A practical rule of thumb is:

- Do not move a workflow from **Assistive → Semi-autonomous** until governance and evaluation are at least “good” (4 out of 5) and there is a clear rollback path.
- Do not move from **Semi-autonomous → Guardrailed autonomy** until the workflow has demonstrated stable performance for a defined period and has passed at least one **formal promotion review**.

5.4. Governance cadence and ownership

For the Ladder to work in practice, two disciplines are crucial:

A regular promotion and demotion cadence

- Every 4–6 weeks, relevant workflows are reviewed: performance against targets, incidents, and changes in context (new regulations, new products, new risks).
- A workflow can move up, stay where it is, or be moved down a level if performance has slipped or risk has increased.
- Decisions are logged and communicated clearly to stakeholders (CX leads, product teams, risk and compliance).

Clear ownership

- **Product / CX owners** articulate the business value, customer impact, and acceptable risk for each workflow.
- **Technology/platform teams** own implementation, observability, and adherence to promotion gates.
- **Risk/compliance** validates that governance and controls are adequate for the proposed level of autonomy.
- **Executives** approve the overall framework and intervene when trade-offs need to be escalated.

The result is that when an executive asks, “What is this AI actually doing today?”, the answer is concrete:

- “In cancellations, we are at **semi-autonomous**: the system can process straightforward cases within policy and hands off everything else.”
- “In billing disputes, we remain at **assistive** while we improve data quality and evaluation.”
- “We have granted **guardrailed autonomy** only in one narrow workflow, with explicit metrics, monitoring, and a kill-switch.”

This level of clarity is what ultimately unlocks board-level trust, and the investment to take CX AI beyond experiments.

6. Channel Automation Frontier: Where to Start and How to Expand

Not all channels are equally ready for automation. The **Channel Automation Frontier** is shaped by feasibility (data, APIs, latency) and risk (impact of errors, expectations of tone and nuance).

6.1. Channel prioritization

First wave: chat and email

High intent clarity, structured context, and mature workflows. Ideal for assistive → semi-autonomous progression.

Second wave: in-app

In-application contexts enable UI confirmations, identity assurance, and policy enforcement. Good candidates for early guarded autonomy when actions are reversible.

Third wave: voice

High potential, but requires robust handling of interruptions, ambiguity, and emotional nuance. Start with tightly bounded intents; pair with human fallback and after-call AI for QA and summarization.

Fourth wave: multi-step transactional flows

These require orchestration across several systems, consistent state management, and non-trivial rollback. They should be addressed once orchestration and evaluation capabilities are mature.

> *"In some journeys it's actually better to route customers straight to a specialist than to push them through another digital flow. AI helps us recognize when to keep them in self-service and when to get them directly to the right person."*
>
> Tjitte Folkertsma, COO, Youwe

Channel Automation Frontier

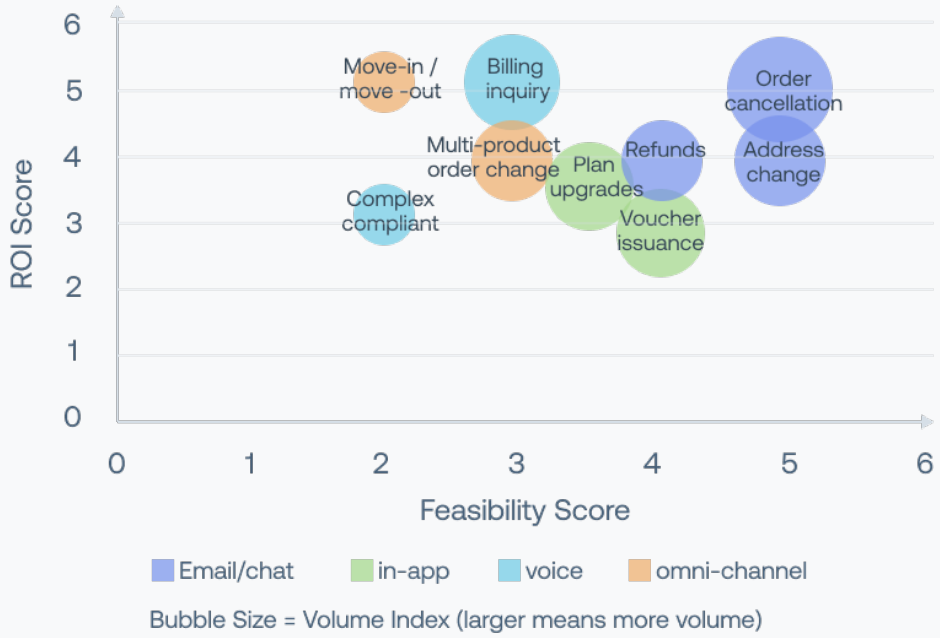

6.2. Guardrails by channel

Across channels, successful organizations share common guardrails:

- **Confidence thresholds** calibrated at the level of individual actions.
- **Dual confirmation** for high-impact changes, especially in voice or transactional contexts.
- **Rich handoff metadata** so agents see the AI's interpretation, evidence, and intended actions.
- **Channel-specific evaluation sets**, updated as behavior and content evolve.

7. Rollout Playbook: From First Workflow to Portfolio

The **Rollout Playbook** translates strategy into a sequence of steps. Most organizations that we interviewed followed a similar pattern:

7.1. Sequencing CX domains

1. **Customer Support:** High volume, relatively standardized intents, and well-defined KPIs. Ideal for initial assistive and semi-autonomous deployments.
2. **Professional Services:** Structured, repeatable tasks (e.g., proposal drafting, status updates, documentation). Attractive for assistive AI and selective autonomy.
3. **Sales:** High stakes; ideal for assistive coaching, note-taking, and follow-up drafting in early phases. Autonomous actions on pricing or commitments should follow only when governance and evaluation are robust.

Rollout Roadmap

Phase	Name	Function	Duration (Days)	Headline Goal
1	Prove the path	Customer Support	15	Baseline metrics and validate feasibility in one high-volume workflow.
2.	Ship assistive	Customer Support	30	Deploy assistive AI into production with A/B and monitoring.
3.	Earn semi-autonomy	Customer Support	30	Allow limited direct actions in low-risk cases; test rollback and kill-switch.
4.	Prepare to scale	Customer Support & Platform	15	Package pattern and hold promotion review; select next workflows.
5.	Extend to Pro Services	Professional Services	30	Apply pattern to structured tasks (proposals, status updates).
6.	Extend to Sales	Sales	30	Deploy assistive coaching and follow-up drafting in sales flows.

7.2. A 90-day implementation guide

A 90-day horizon is enough to prove the path and establish patterns.

Days 0–15: Prove the path

- Select one workflow with meaningful volume and clearly defined success metrics.
- Baseline containment, AHT, CSAT, FCR, and cost per contact.
- Build a shadow-mode prototype and an offline evaluation set.

Days 16–45: Ship assistive

- Deploy assistive capabilities with A/B testing.
- Implement policy filters and logging.
- Begin weekly regression testing and error analysis.

Days 46–75: Earn semi-autonomy

- Introduce semi-autonomous actions in narrow scopes (act-by-exception).
- Promote to act-by-default only where performance and error rates meet thresholds.
- Test rollback and kill-switch mechanisms.

Days 76–90: Prepare to scale

- Package a reusable pattern (orchestration template, policy rules, evaluation suite, dashboards).
- Conduct a Promotion Review with CX, IT, and Legal.
- Identify the next two workflows to apply the pattern.

> *“This technology moves so fast that it’s more realistic to start with a proof of concept and learn from that.”*
>
> Pablo Diez Garcia, Process Transformation & Innovation Manager, Generali

7.3. What “good” looks like after 12–18 months

- A decoupled agent layer with explicit policy and evaluation services.
- Governance and evaluation ≥4/5; orchestration and data readiness ≥3.5/5.
- A portfolio of 6–10 workflows in assistive or semi-autonomous modes, with 2–3 in guardrailed autonomy.
- Demonstrable improvements in containment and AHT, with stable or improved CSAT and FCR.

Risk & Governance Checklist:

- ☐ **Data & Privacy:** classification, masking/tokenization, retention and sovereignty documented.
- ☐ **Governance:** action rights aligned with Agentic Ladder levels; clear HITL policies; audit trails for prompts and actions.
- ☐ **Evaluation:** offline test sets with thresholds; online safety monitors; regression gates.
- ☐ **Operations:** incident response playbook; kill-switch and rollback drills; weekly performance scorecards.
- ☐ **Value:** KPIs for containment, AHT, FCR, CSAT, and cost per contact, with explicit target directions.

8. Closing Note

The interviews underlying this whitepaper convey a consistent message: the decisive factor in CX AI is not who deploys the most models, but who builds the most disciplined operating system around them. Autonomy follows maturity, not ambition.

Organizations that succeed do three things well:

1. They frame AI around workflows, not generic capabilities.
2. They connect deployment decisions to maturity gates and sourcing choices, not to enthusiasm or fear of missing out.
3. They build and reuse patterns, for evaluation, policy, and orchestration, so that each successful use case lowers the cost and risk of the next.

In that sense, AI in CX is less a technological race and more a test of institutional learning.

> *"AI is definitely on our agenda, and I'm convinced it will only become more and more important. The real question now isn't if we use it, it's how we do it in the right way."*
>
> Claude Pannier, CFO, SPIE

About Typewise

Typewise is an enterprise-grade AI customer service platform that fundamentally optimizes and automates customer service with the help of custom AI Agents and an intelligent Agent Assistant. Numerous Fortune 500 companies, including Unilever and DPD, have reduced effort by 50% or more using Typewise's cutting-edge AI technology.

Typewise is backed by Y Combinator (S22) and has repeatedly been recognized as a "High Performer" in AI-driven automation on G2.

David Eberle – CEO & Co-founder, Typewise

David Eberle has over a decade of experience in tech strategy and has led multi-million-dollar transformation projects for Fortune 500 companies. A global entrepreneur, he has lived in eight countries and speaks six languages.

Janis Berneker – CTO & Co-founder, Typewise

Janis Berneker leads AI and technology development at Typewise. A lifelong AI innovator, he previously built data science models for Swiss national TV (SRF) and has extensive experience in large-scale data analytics and automation.

Contact

David Eberle – CEO & Co-founder | david.eberle@typewise.app

Appendix A – CX AI Maturity OS

The CX AI Maturity OS scores **each workflow**, not the organization overall, on **six dimensions** from **0 to 5**.

Scores are interpreted as: 0 = not present, 3 = good/repeatable, 5 = leading practice.

1. Strategy – objectives beyond experimentation

- **0** – Ad hoc pilots; no clear CX or cost goals; no named owner.
- **3** – Documented CX AI goals linked to KPIs; named executive sponsor; basic roadmap.
- **5** – CX AI embedded in business strategy; portfolio of workflows managed with regular exec reviews.

2. Data readiness – knowledge and data for AI

- **0** – Knowledge scattered across tools, no trusted source of truth.
- **3** – Central knowledge base for key CX domains; owners defined; basic update process.
- **5** – Systematic knowledge management (taxonomy, lifecycle, quality checks) used by agents and AI.

3. Governance & risk – policies, ownership, auditability

- **0** – Few written policies; risk decisions are ad hoc; minimal logging.
- **3** – Policies exist for sensitive domains; cross-functional review for major AI changes; basic audit logs.
- **5** – Formal governance framework; clear RACI; traceable decisions that satisfy audit/regulator expectations.

4. Stack orchestration – directing and constraining actions

- **0** – Isolated experiments; little or no API access to core systems.
- **3** – API coverage for primary workflows; defined toolset with guardrails and basic observability.
- **5** – Dedicated agent/orchestration layer across channels; reusable actions; rich traces and metrics.

5. Talent & enablement – people and change

- **0** – Little awareness or training; AI usage is optional and informal.
- **3** – Structured training for agents and supervisors; feedback loops into product/tech.
- **5** – AI literacy part of standard onboarding; clear expectations and monitoring; defined "AI champion" roles.

6. Value realization – evidence of impact and ability to scale

- **0** – No baselines; success described anecdotally.
- **3** – Baselines and targets for key workflows; measured improvements (e.g., AHT, containment, CSAT, cost/contact).
- **5** – Portfolio view of workflows; consistent KPI framework; value reported regularly to executives and used to prioritize investment.

Usage: Workflows are scored on these six dimensions. Promotion up the **Agentic Ladder** (assistive → semi-autonomous → guardrailed autonomy) is granted only when agreed thresholds are met (e.g., governance & risk ≥4, stack orchestration ≥3.5, value realization ≥3).

Appendix B – Methodology & Interview Panel

B.1 Methodology

Sample: 36 structured interviews with C-level and senior leaders across Europe.

Roles: Chief Customer Officer, Chief Digital / Technology Officer, Chief Experience Officer, Head of Customer Service / Contact Centre, and equivalents.

Industries (grouped):

- **Tech & Software** – software/SaaS, IT services, internet
- **Consumer Services & Retail** – retail / e-commerce, gambling & casinos, travel & hospitality, other B2C services
- **Commercial Services** – automotive, facility management & real estate, business supplies, logistics / B2B services
- **Regulated & Infrastructure** – insurance, public sector, rail, utilities

Approach:

- Transcripts coded against a normalized schema (adoption stage, agentic posture, sourcing posture, use cases, blockers/enablers, quotes).
- Six-dimension maturity scores (0–5) assigned per workflow using the CX AI Maturity OS (Appendix A).
- Records joined with firmographics: industry, employee band (0–199, 200–999, 1k–4.9k, 5k–9k, 10k+), country/region.
- Aggregations built for charts (e.g., adoption by industry, sourcing by size, maturity by cluster).

Key metrics used:

- Containment rate, AHT (Average Handling Time), CSAT, FCR (First Contact Resolution), cost per contact.

Limitations:

- Qualitative panel (36 interviews), not a statistically representative survey.
- Participants skew toward organizations already engaged with CX AI.

- The market is evolving rapidly; frameworks are designed to be robust to change, but specific tools will shift.

B.2 Interview Panel Overview

#	Name	Title	Company	Industry	Employee band	HQ Country
1	Martin Sebelius	CEO Accedo Video Solutions	Accedo	Software	501-1k	Sweden
2	Olja Draskovic	Head of Customer Experience	AdmiralBet	Gambling and Casinos	1k-5k	Serbia
3	Adriano Ciarletti	CEO	Bricocenter	Retail	1k-5k	Italy
4	N.N.	Head of Innovation Lab	Government Agency	Public Sector	5k-10k	Germany
5	Peter Schneck	CEO	CENIT	IT Services & Consulting	501-1k	Germany
6	Francesca Meriggi	Group CIO	Engineering Group	IT Services & Consulting	10k+	Italy
7	Hans-Kaspar Scherrer	CEO	Eniwa AG	Utilities	201-500	Switzerland
8	Pablo Diez Garcia	Process Transformation & Innovation Manager	Generali	Insurance	10k+	Spain
9	Bard Lund	Country Managing Director	Hitachi Energy	Utilities	10k+	Norway
10	N.N.	CEO	Car Dealer Group	Automotive	501-1k	Germany
11	Michal Paschalis	CEO	IAI Group	IT Services & Consulting	201-500	Poland
12	Tiago Goncalves	CEO	InnoWave	IT Services & Consulting	201-500	Spain
13	Hueseyin Dogan	COO	IONOS	IT Services & Consulting	1k-5k	Germany
14	Robert Vermin	CFO	Just Brands	Retail	201-500	Netherlands
15	Carsten Nørrevang	CFO	Kirppu	Retail	201-500	Denmark
16	Beat Gisin	CFO	Koelliker Büro-automation	Business Supplies	51-200	Switzerland
17	Ross Abbate	CEO	Macro	Facilities Services	501-1k	United Kingdom
18	Martijn Willems	CFO	Mantel	Retail	501-1k	Netherlands
19	Jon-Gunnar Aasen	CEO	Miles AS	IT Services & Consulting	201-500	Norway

20	Ana Combalia	CEO	Natura Italia	Retail	51-200	Italy
21	Emilija Frew	Head of Tech	Pierce Group	Retail	201-500	Sweden
22	Pavlo Antonov	CTO	SelectSpecs	Retail	51-200	United Kingdom
23	Frank Robben	CEO	Smals	IT Services & Consulting	1k-5k	Belgium
24	Thibault Milan	Head of Innovation	Smile	IT Services & Consulting	1k-5k	France
25	Claude Pannier	CFO	SPIE Belgium	Facilities Services	1k-5k	Belgium
26	Claus Nolte	Head of IT Governance & Strategy	Stadtwerke Düsseldorf	Utilities	1k-5k	Germany
27	Davide Conigliaro	CEO Humatics	SYS-DAT Group	IT Services & Consulting	501-1k	Italy
28	Ida La Spisa	CIO	Telia	Telco	10k+	Sweden
29	Steve Finlan	CEO	The Wine Society	Retail	201-500	United Kingdom
30	Antonio Tresca	Head of Customer Service and Direct Sales	Trenitalia	Rail Transport	10k+	Italy
31	Charlotte Moerch	Director, Head of Customer Relations	Tryg	Insurance	5k-10k	Denmark
32	N.N.	Head of Information Systems	Global Automotive Company	Logistics	1k-5k	Germany
33	Karsten Rech	CIO	Vonovia	Real Estate	10k+	Germany
34	Tjitte Folkertsma	CIO	Youwe	IT Services & Consulting	201-500	Netherlands
35	Sylwia Bischof	Head of Customer Service DTC	Zattoo	IT Services & Consulting	201-500	Germany
36	Soren Meyer	CEO	ZeroNorth	Software	501-1k	Denmark