AIGW | Global AI Governance & Workforce Transformation Policy Observatory

AIGW ENTERPRISE WORKFORCE TRANSITION · DEEP DIVE 16

Beyond the 700-Agent Headline

Klarna and the Workforce Economics of AI-First Customer Service

Evidence reviewed through 11 November 2026

ENTERPRISE CASE STUDY

Executive summary

The story most people remember about Klarna is simple: an AI customer-service assistant did the work of 700 people; then the company had to bring humans back. It is a memorable narrative, but the two halves are not a faithful account of either workforce measurement or organizational change.

In February 2024, Klarna said its OpenAI-powered assistant handled 2.3 million conversations in its first month, approximately two-thirds of service chats. It described the work as equivalent to 700 full-time agents, claimed customer satisfaction on par with human service, and projected a $40 million profit improvement that year. Those figures were company-reported initial indicators, not a count of employees dismissed or independently audited profit. S01

The meaning of the 700-person claim is unusually important. In a 2024 interview, CEO Sebastian Siemiatkowski explained that the company used external customer-service providers and that the calculation reflected reduced demand for their agents; Klarna did not itself employ those outsourced service agents. Later Securities and Exchange Commission disclosures identified the method as an estimate based on reductions in average monthly chat and telephone conversations handled by human agents. This is evidence about a **purchased service-capacity requirement**, not an employee-level ledger, a validated labor displacement estimate or a census of workers' subsequent outcomes. S02 S03

ENTERPRISE CASE STUDY

In May 2025, the CEO said the company had overemphasized cost and was developing a way to ensure that customers could reach human representatives. Yet reporting based on a company spokesperson described the immediate experiment as **two highly experienced agents in a small pilot**, not mass rehiring. Klarna maintained its investments in automation and subsequently disclosed that the assistant handled 80% of customer-service chats in 2025, represented more than 850 full-time-agent equivalents under its estimation method and generated approximately $59 million in claimed 2025 savings. Its annual report calls the result a *dual-track* approach: scaled AI service with a continuing human option. The widely used phrase 'AI reversal' obscures this operational coexistence. S04 S05 S07

The company also reported its own employee count falling from approximately 4,352 at the end of 2023 to 3,422 at the end of 2024 and 2,831 at the end of 2025. Its annual report explicitly connects its overall headcount strategy to leveraging AI and operational efficiency. But those stock figures belong to Klarna's corporate workforce; the original 700-agent claim belongs to contracted service capacity. The public sources do not provide the inflows, voluntary exits, non-backfill, restructuring, provider contract volumes or individual destinations needed to reconcile them. No arithmetic transformation can repair that evidentiary gap. S07

The result is a revealing case for the design of AI-native service organizations. Routine demand can be handled at scale while customers, complex disputes and trust-sensitive situations continue to require a credible human path. A firm may genuinely reduce service cost and simultaneously conclude that a cost-only design is insufficient. Yet if it does not measure the people performing outsourced work, the human service workforce can become invisible in official corporate headcount reporting.

ENTERPRISE CASE STUDY

This Deep Dive examines the sequence, the service work, the limits of the metrics, the human-support pilot and the unresolved workforce and governance questions. It compares Klarna with JD.com's longer customer-service automation experience in China and interprets the difference through China's October 2026 employment-friendly AI directive. These comparisons do not establish which firm has achieved superior worker outcomes. The report's institutional conclusion is narrower and more useful: **AI-first customer service needs a service-quality architecture, a worker-flow accounting system that includes contractors, and enforceable human escalation and responsibility—not simply an automation rate.**

1. Research question, episode and method

Question: How did Klarna's AI assistant change the demand for service labor, the organization of human support and the evidence required to demonstrate a successful workforce transition?

We distinguish six units of analysis: (1) conversations handled by AI; (2) time and quality of service; (3) estimated equivalent external-agent capacity; (4) contractor and vendor workforce flows; (5) Klarna's own corporate employee stock and movements; and (6) customer and worker outcomes. A number at one level must not be silently transferred to another. The source set combines company announcements and filings, an executive interview that identifies the outsourced-worker denominator, company-spokesperson responses, independent reporting, external industry surveys and government policy. This is a documentary case using public evidence; it does not include customer logs, procurement contracts, independent satisfaction samples or private employee records.

The observed episode runs from the 2024 launch and customer-service transformation through 2025 efforts to strengthen human access and the latest 2026 operating disclosures available by 11 November. The company's scheduled Q3 2026 earnings release on **17 November 2026** lies outside this review window and must not be incorporated retrospectively as if its results were already known. S09

This report does not assert that AI is the only cause of Klarna's cost changes, hiring practices or financial results. Expansion, product mix, channel migration, demand growth, contracted service arrangements and other restructuring may all affect these measures.

MEASUREMENT BOUNDARIES

Three different units in the public record

FIGURE 1

Distinct measures, not equivalent

Service activity, purchased capacity and employees

Measure	Reported value and scope	What it does not establish
AI conversations	2.3 million in the first month, 2024.	Not unique customers or confirmed final resolutions.
External capacity	More than 850 agent-equivalents estimated for 2025.	An estimate of purchased capacity, not verified layoffs or distinct workers.
Direct employees	Approximately 2,831 full-time company employees at year-end 2025.	A company-wide employee stock, not a count of AI-caused service exits.

Source: Klarna disclosures (references S01, S03 and S07) [1] [3] [7]

Note: These are different populations, periods and measurement units. Company-reported claims are not independently audited merely because they appear in a filing.

ENTERPRISE CASE STUDY

2. The service operating model before the headline

Customer service in payments and commerce is not one task. It includes requests about purchases, account access, returns, refunds, payment arrangements, disputes, language barriers and exceptions that cross organizational boundaries. Some requests are repetitive and information-rich. Others require investigation, human judgment, empathy, authority to remedy harm or coordination with merchants and financial-service partners. Treating this entire system as a single queue masks the different failure and escalation costs associated with each interaction.

Klarna's original organization relied heavily on external service providers. The CEO publicly explained that Klarna did not directly employ the outsourced service agents behind the famous 700-agent comparison. This distinction matters for governance and workforce measurement. Automation may first reduce contracted volumes, alter vendor staffing or change what providers are paid to perform. Such changes can precede, coexist with or occur without a visible change in the purchaser's reported headcount. S02

A complete baseline would include the volume of eligible contacts by issue class and channel; contractor FTEs by supplier and geography; staffing by tier and language; case complexity; first-contact resolution; follow-up effort; repeat contacts; fraud and financial-risk escalations; customer access to people; and average cost per resolved case. The public record offers selected parts of this baseline but not the whole denominator. In particular, an estimate based on reduced human-handled conversations does not automatically reveal the number of supplier employees who lost jobs, were reassigned to other clients or remained employed through reduced hours.

ENTERPRISE CASE STUDY

3. 2024: a large-scale rollout with unusually ambitious claims

On 27 February 2024, Klarna announced that its AI assistant had participated in 2.3 million conversations after a month of deployment and handled roughly two-thirds of customer-service chats. It reported a similar customer satisfaction score to human agents, a 25% decline in repeat inquiries, and resolutions in less than two minutes rather than an earlier human-service benchmark of 11 minutes. The initial statement projected roughly $40 million of profit improvement in 2024. It also described service in 23 markets and over 35 languages. The service-scale claims establish that the system was not a laboratory experiment. They do not independently validate the comparative samples, control for case mix or establish the projected earnings impact. S01

Later filing language refines several headline measures. Klarna reported approximately **$39 million of cost savings for 2024**, rather than presenting the original projected $40 million profit improvement as an audited realized outcome. The filing says the over-700 agent-equivalent estimate was based on the average monthly reduction in conversations handled by human agents after AI launch. It also reports that 2024 AI-handled resolutions averaged roughly two minutes compared with a 12-minute average for human agents, under the later company's reporting period and definitions. The change from an 11-minute early comparison to 12 minutes in subsequent filings is not a reason to construct a time series without the underlying measurement design. S03 S07

There are at least four possible sources of misinterpretation. First, conversations are not unique customers. Second, a chat may be completed or routed in ways that do not mean the same thing as fully resolved financial issues. Third, a short average time may reflect the selection of more straightforward requests for AI. Fourth, the estimate of displaced *service demand* does not identify the individual service workers affected by vendor contracts. These are not reasons to dismiss the rollout; they are reasons to state exactly what it demonstrates.

ENTERPRISE CASE STUDY

4. The 700-agent number: service capacity, not 700 layoffs

The case's most frequently repeated quantitative claim is also its most susceptible to category error. The statement that an AI assistant performs 'the work of 700 full-time agents' sounds like a count of workers. In fact, the company expressly described the reduction in human conversation demand through outsourced service providers. Later regulatory documents define it as a counterfactual labor requirement inferred from reductions in human-handled chat and telephone traffic, not a direct measurement of 700 named positions eliminated. S02 S03

For the employer purchasing services, lower required external agent capacity can be a real economic benefit. For an outsourced provider, the same decision might mean shorter shifts, fewer agency hours, reallocation to another contract, lower hiring, attrition or formal redundancies. The public evidence does not distinguish those possibilities. Nor does it show how much work was shifted into quality assurance, AI operations, dispute resolution, human escalation and system maintenance.

Three denominators should remain separate on every page of analysis: **AI-handled service conversations**, **estimated purchased human-agent capacity**, and **Klarna's own corporate employees**. Their numerical values have no common unit. A workforce outcome requires actual workforce flows, not a multiplication of 'time saved' by an assumed wage or hours-per-worker factor.

We therefore reject two statements as unsupported: that Klarna laid off 700 of its own customer-service employees because of the AI assistant, and that each agent-equivalent capacity reduction corresponds one-for-one to a worker dismissed by an outsourcer. Both exceed the observed evidence.

ENTERPRISE CASE STUDY

5. 2025: the human-service correction—and its limits

In May 2025, Bloomberg reported that the CEO considered the firm's focus on cost reduction to have gone too far and was developing arrangements under which consumers could always reach a real person. The narrative quickly became a reversal: the AI company was 'bringing humans back.' A closer reading of contemporaneous evidence gives a more interesting picture. Company spokespeople told independent reporters that the human-access strategy was a refinement and that AI remained central. A company spokesperson quoted by CIO described **two highly experienced agents hired for a small pilot**, while rejecting the idea that the firm was generally rehiring former departing staff. S04 S05 S06

The CEO also spoke publicly in June 2025 about offering more human service, including premium or high-touch interactions. These remarks signal an intended redesign of the service mix; they do not establish the size of a newly employed human support organization or the ultimate contractual arrangement. S18

The organizational lesson is not that AI's scale benefits were fictional. It is that **average cost and handling speed can become over-dominant performance criteria**, while difficult cases and customer preferences require distinct service capabilities. A two-agent pilot establishes experimentation, not proof that Klarna created a durable new internal career ladder. Neither does a commitment to maintain human access establish consistent waiting times, case-handling authority or error-repair rights for every customer.

The difference between the two narratives is analytically consequential. A dramatic 'AI reversal' headline suggests technological failure; an 'AI plus humans' corporate slogan suggests the problem has already been solved. Both are stronger than the public evidence. The observed transition is continued scaled automation alongside a more explicit human-service boundary.

ENTERPRISE CASE STUDY

6. 2025 annual reporting: AI volume and the human option coexist

Klarna's 2025 Form 20-F is the strongest primary source on the subsequent operating model. It reports that the assistant handled **80% of customer-service chats in the year ended 31 December 2025**, with approximately **31 million cumulative conversations since launch**. Its estimate of service-capacity equivalence rose to **more than 850 full-time agents**, using 2025 reductions in human-handled chat and telephone traffic. It estimates **$59 million in customer-service cost savings in 2025**. In the same section it states that customers retain the option to speak to human representatives. S07

These are not mutually exclusive claims. A service system can automate a large share of chats while preserving a small, expensive and strategically significant human pathway. The marginal value of humans may be highest in cases involving complaints, vulnerable customers, exceptions, misunderstandings, fraud concerns or recovery of trust. The filing says satisfaction was on par with human agents according to **internal** surveys; it does not publish enough distributional evidence to conclude the quality of all complex interactions or verify how often a human option was exercised.

The company also attributes operating leverage to scale and multiple efficiency initiatives, with AI a named contributor. For analytical purposes, an estimated $59 million service-cost saving is a **company-reported cost attribution**, not a third-party causal estimate of net economic value. It should be assessed against model costs, oversight, supplier changes, contact volume and customer outcomes. The 2024 $39 million and 2025 $59 million figures are for different periods; they should not be summed into a made-up lifetime net profit unless their treatment and counterfactuals are independently reconciled.

ENTERPRISE CASE STUDY

7. Workforce accounting: two labor markets, two different records

Klarna's annual report reports **approximately 4,352 direct full-time employees at year-end 2023, 3,422 at year-end 2024, and 2,831 at year-end 2025**. Management explicitly describes a strategic decision to reduce headcount and seek efficiencies through AI. The observation is important and should not be downplayed. But the reported annual counts are total corporate headcount stocks, not a distribution of departures by function, route or cause. S07

The arithmetic difference between the 2023 and 2025 stocks is 1,521 fewer direct employees. That is a net change, not evidence of 1,521 layoffs. It could reflect combinations of voluntary departures, hiring restraint, non-backfill, restructuring, entity changes and involuntary exits. Even if the company explicitly associates the strategy with AI, its public filings do not isolate the marginal causal contribution of the customer-service assistant versus other AI use or wider efficiency programs.

The external service workforce is even less visible. The AI agent-equivalent claims describe contracted human capacity, while vendor employment records remain absent. An accurate transition study would request supplier-level baseline headcount and paid hours, contracts gained or lost, redeployment rates, employment status after 3/6/12 months, changes in wages and conditions, and the staffing of human escalation, quality assurance and knowledge operations. It would also ask whether suppliers absorbed demand on other clients and how much variation exists across languages and markets.

This is why AIGW treats workforce transition as a **stock-and-flow** question rather than a single headcount trend. If a company reduces recruitment, another lets short-term contracts lapse, and a third redeploys people to specialized care, those decisions have different effects on livelihoods, experience accumulation and accountability even if each reports the same 'agent-equivalent' figure. This case does not allow us to resolve those differences.

ENTERPRISE CASE STUDY

8. Service quality is a workflow design problem, not one satisfaction score

The fundamental organizational problem in an AI-first service environment is where the automation boundary sits. Routine requests may be answered immediately and accurately, while disputes can require legal context, financial verification and discretion. A transparent service architecture therefore needs at least five measurable stages: initial intent classification, AI handling, exception detection, human transfer, and verified resolution. It should specify the customer-facing handoff and the authority of the human recipient. Independent Reuters reporting in December 2025 likewise described Klarna's continued reliance on humans for complex cases while AI handled simpler interactions. S17

Klarna's filings describe fast resolution averages, internal satisfaction surveys and continuing human access. They do not release enough underlying data to reconstruct these steps end-to-end. We cannot determine whether the 80% handled volume includes transferred contacts, whether escalated cases have different quality outcomes, or the proportion of customers who requested a person and actually reached one within an acceptable service level. The appropriate response is to define these as **evidence gaps** rather than insinuate a failure for which evidence is absent. S07

A useful managerial measurement set would include: correct final resolution by issue type; customer-requested and system-triggered escalation; time to a qualified human; repeated contact within 7 and 30 days; reversal and refund accuracy; complaint recovery; accessibility by language; human overrides; audit trails for consequential actions; and the full labor and technology cost per *durably resolved* case. An average chat-handling duration alone can reward premature closure if follow-up work is invisible.

The role of a human service professional may also become harder, not easier. If AI resolves simple queries, the remaining distribution may contain higher emotional intensity, greater ambiguity, stronger financial stakes and more cross-team coordination. That raises questions about staffing ratios, breaks, specialist support, escalation authority, training, compensation and burnout. We have no public worker-level dataset to determine whether Klarna's support providers experienced these outcomes; the research agenda should not convert this credible mechanism into an observed effect.

ENTERPRISE CASE STUDY

9. The business ledger cannot be reduced to 'revenue per employee'

Klarna reported that annual revenue per employee increased from approximately $344,000 in 2022 to $1.24 million in 2025. That is a striking operating-leverage indicator, but it combines changes in revenue, business scale, products and corporate headcount. It should not be used to derive an estimate of service-agent productivity or to attribute the entire gain to the chatbot. S07

The 2026 Q2 report offers a useful caution. Klarna reported revenue growth of 27% year over year, while adjusted customer-service and operations expense increased from $48 million to $57 million. The company also reported strong growth in users, merchants and business volume. These figures show that cost is not necessarily a monotonic function of automation penetration: service volume, operating reach, case mix and investment can rise together. They do not prove that AI worsened service economics, nor that the human-service pilot caused the expense increase. S08

A stronger financial model would segment cost per verified resolution, lifetime customer value and retention, disputes, vendor expenses, model inference and infrastructure, AI maintenance, human quality assurance and resolution of errors. Its denominator should reflect demand and outcome, not just the size of the corporate payroll. Management may rightly decide that maintaining a human option is worth the cost when the expected loss of trust or customer exit would be greater. That proposition is strategically plausible but cannot be quantified from currently published material.

ENTERPRISE CASE STUDY

10. A wider service-industry test: the 'cut fast, rehire later' risk

An external survey helps interpret why the human-service boundary matters beyond Klarna. In November 2026, Gartner reported a survey of 275 service and support leaders conducted in July–August: 21% said their organizations had reduced human-agent headcount because of AI, 7% reported rehiring following previous reductions, and 11% planned to rehire within 18 months. Those are **responses among surveyed leaders**, not percentages of workers, Klarna's own experience or causal estimates of industrywide employment. The survey also describes firms moving people into other functions and hiring specialists for AI initiatives. S10

A related survey release reports that only 4% of customer-service organizations had realized their expected return on AI initiatives. That statistic reflects self-reported expectations within the sampled organizations, not an audited failure rate of AI projects. It reinforces the need for fuller service outcome accounting while leaving room for some organizations to achieve genuine savings. Klarna's filings supply their own company-reported operating gains; neither source validates or invalidates the other. S11

The plausible failure mode is organizational rather than simply technological: leaders adjust a staffing plan against forecast automation capacity before the system has demonstrated reliable performance in high-stakes or exceptional cases. Recruitment, training and knowledge restoration then become expensive. This is an analytical interpretation of Gartner's survey and the Klarna sequence; it is not evidence that Klarna dismissed and later rehired exactly the same workers. S10

ENTERPRISE CASE STUDY

11. China's service-work transition: JD.com as a mechanism comparator

China provides a useful comparison because customer-service automation has developed within large platforms for years, not only since the recent wave of large language models. JD.com reported in its 2025 results that its AI service system handled more than **4.2 billion customer inquiries during the 2025 11.11 Grand Promotion** across the group's businesses. This is an enormous scale indicator, but it is not directly comparable with Klarna's 31 million *cumulative* assistant conversations: the periods, contact units, covered businesses, channels and definitions differ. The correct use is to show that both organizations operate large AI service systems, not to rank their replacement rates or performance from dissimilar counts. S12

AIGW's separate JD.com enterprise investigation documents a qualitatively different kind of workforce trace: a customer-service professional with nearly a decade of experience moving into AI training and operations, with work such as diagnosing response failures, modifying knowledge content, configuring agents and maintaining the service system. That worker-level example is stronger evidence of a possible destination role than an anonymous 'agent-equivalent' estimate. Yet it does not prove how many employees reached those roles, whether they earned more, or whether the role is durable as AI systems improve. The underlying primary reports are traced in the AIGW JD dossier and should not be treated as independent replication simply because AIGW has published a synthesis. S13

Together, Klarna and JD.com suggest three distinct work movements: routine response work goes to AI; complex and trust-sensitive service remains with specialized humans; and some human expertise moves into configuring, evaluating or supervising the AI operating system itself. The balance and distribution differ by institution and are not identified with equal evidentiary strength. The central international question is no longer only *how many contacts are automated*, but *how service organizations account for the human work that changes or remains*.

ENTERPRISE CASE STUDY

12. China's employment-friendly AI policy: a different institutional layer

The Chinese State Council General Office issued **Guobanfa [2026] No. 29**, adopted on 16 October and published on 23 October 2026, concerning adaptation to AI and construction of an employment-friendly development model. The directive addresses job creation, reduction of displacement risk, worker capabilities, social protection and employment-risk monitoring. It calls for institutional responses alongside technical adoption, including attention to human–machine collaboration and worker transition. These are prospective policy commitments and design requirements; they are not a causal explanation of JD.com's previous company practices and they do not establish that China has solved service-sector transitions. S14 S15

For an outsourced AI service operation, the directive highlights questions that are absent from headline productivity numbers: who is responsible for assessing displacement risk when labor is supplied through a contractor; which learning opportunities and transfer routes are available; how nonstandard and flexible workers obtain protection; and which data should be reported to public employment institutions. Applying those questions to Klarna is an AIGW comparative inference, not an assertion of Chinese policy jurisdiction over a Swedish fintech.

A complementary European perspective concerns transparency to users interacting with AI systems. Article 50 of the EU AI Act contains requirements on informing people when they interact directly with relevant AI systems, subject to specified conditions and exceptions. This is a relevant governance design consideration for chat-based service. It does not by itself establish a legal right in every circumstance to a human agent, or prove Klarna's compliance status. This report makes no company-specific legal compliance determination. S16

The comparative architecture is therefore three-layered: enterprise service design, workforce transition infrastructure and customer-facing accountability. Policy may reshape incentives and rights, but its effects must be measured rather than presumed. The distinctive China contribution is an institutional lens for monitoring employment risks across suppliers and supporting transitions, not a claim that a regulation automatically generates good jobs.

ENTERPRISE CASE STUDY

13. A practical framework: the human-service boundary

The case suggests a management method for designing an AI-first service system without using headcount cuts as the starting assumption.

Gate 1 — Work eligibility. Segment tasks by intent, regulatory and financial consequence, complexity, ambiguity and customer vulnerability. Automation-rate targets should be attached to bounded groups, not a single undifferentiated service category.

Gate 2 — Verified resolution. Measure quality beyond answer speed: correct resolution, repeat contacts, appeals, error rates, dispute outcomes and customer retention. Audit difficult requests separately rather than pooling them with routine queries.

Gate 3 — Human access and authority. Give customers a usable escalation path; measure transfer success and wait times; specify what representatives can correct, refund, override or elevate; and assign accountability for AI-origin errors. A human response box without remedial authority is not meaningful oversight.

Gate 4 — Workforce accounting across boundaries. Track internal and outsourced headcount, paid hours, hiring, departures, reduced recruitment, non-backfill, redeployment and new AI operations roles. Obtain supplier-level data where contracts permit. Record worker outcomes separately from estimated 'agent-equivalent capacity.'

Gate 5 — Career and learning architecture. If routine tickets formerly trained novice staff, create deliberate exposure to supervised exceptions, judgment exercises, customer context, quality evaluations and AI-system maintenance. Track competence and access to actual assignments rather than counting only training hours.

Gate 6 — Business outcome and reversibility. Review net economics of automation, model and oversight cost, quality failures and reinvestment. Establish thresholds for adding human capacity or rolling back an AI workflow before service damage accumulates. Design reversibility ex ante rather than treating re-hiring as a public-relations failure.

These six gates are **AIGW analytical recommendations**, derived from evidence and gaps in the Klarna episode and external comparisons; they are not a tested proprietary maturity score or evidence that Klarna has implemented each control. A controlled pilot should validate the measurement burden, reliability and usefulness before the framework is commercialized.

ENTERPRISE CASE STUDY

14. What the evidence does not settle

Klarna provides unusually strong public evidence of service scale, an explicit capacity-estimation method, material company-reported cost savings and a documented decision to emphasize human service availability. The evidence is weaker on the service workforce as a population of identifiable people. The main unresolved questions are consequential:

Customer outcomes: What share of customers requested human support, reached it and resolved their problems, stratified by issue type, language and vulnerability? Are AI and human satisfaction measurements case-mix adjusted? What happens after apparently completed chats?

Vendor workforce: How many individual outsourced agents changed employment status or working hours? How many were redeployed to other accounts? Did pay, contract security, career development or working conditions change?

Internal headcount: What were the hires, voluntary departures, non-backfilled roles, involuntary exits and internal transfers underlying each year's company employment stock? What share of these flows was linked to the service assistant rather than other initiatives?

Human-service pilot: Did the two-person pilot lead to sustained recruiting, a different supplier mix, specialized positions or a scaled hybrid queue? What were the costs, service outcomes and employee conditions?

Causal finance: How would service cost, resolution quality, growth and financial results have evolved without the AI assistant? Are model, human oversight and incident-resolution costs included in claimed savings?

A complete answer would require disaggregated operating data and worker records that are not public. The next announced financial reporting event is 17 November 2026, beyond this report's cut-off. Future filings may change the assessment but do not retroactively close today's gaps. S09

ENTERPRISE CASE STUDY

15. Conclusion — capacity is not the service organization

Klarna demonstrates that an AI assistant can become a large-scale service operator while the organization simultaneously recognizes the need for people. Its 2025 report records AI handling 80% of customer-service chats, more than 850 agent-equivalent estimated capacity and about $59 million of claimed annual savings; it also describes continuing human service. These coexist because service is not merely a volume of conversations. It is an institutional commitment to reach a correct outcome when a customer confronts uncertainty, an error, an exception or a consequential decision. S07

The company's headcount strategy deserves scrutiny, but its corporate headcount cannot be substituted for the labor conditions of outsourced service workers. A serious workforce study must follow the service contract, the work unit and the worker—not just the conversational interface. The '700 agents' headline was useful for communicating the scale of an operational change, but it is not a workforce transition ledger.

For other enterprises, the durable lesson is not 'automate less' or 'put humans back everywhere.' It is to design the **human-service boundary** before converting technical capability into staffing assumptions: reliable resolution, meaningful escalation, supplier-inclusive workforce accounting, staff competence and a way to restore capacity if the system fails. Comparing Klarna and JD.com shows that companies can reorganize service and AI operations in different ways; China's 2026 policy adds a wider institutional question about who provides the safety and transition infrastructure when the unit of employment is outside the technology adopter.

The decisive question for the next phase of AI adoption is not simply how many customer conversations a model can handle. It is **who is responsible for the difficult cases, and what happens to the people and institutions that previously carried that responsibility**.

SOURCES

Original sources and further reading

The linked sources below are cited in the text with their original publisher URL. Company-disclosed operational figures are not independently audited solely by virtue of appearing in a regulatory filing.

- **[S01] Klarna Bank AB (2024-02-27).** Klarna AI assistant handles two-thirds of customer-service chats in first month.

- **[S02] CBS News (2024).** Interview with Klarna CEO on AI assistant and outsourced agents.

- **[S03] Klarna Group plc / US SEC (2025).** 2025 amended registration statement, initial service and workforce metrics.

- **[S04] Bloomberg News (2025-05-08).** Klarna slows AI-driven job cuts with call for real people.

- **[S05] CIO / Paul Barker (2025-05-12).** Mixed messages from Klarna about plans for more AI, fewer humans.

- **[S06] Customer Experience Dive (2025-05-09).** Klarna changes AI tune and again recruits humans.

- **[S07] Klarna Group plc / US SEC (2026-02-26).** 2025 Form 20-F Annual Report (filed February 2026).

- **[S08] Klarna Group plc / US SEC (2026-08-18).** Klarna second-quarter 2026 earnings release.

- **[S09] Klarna Group plc (2026-10-22).** Klarna Q3 2026 earnings scheduled for November 17.

SOURCES

- **[S10] Gartner (2026-11-05).** Survey finds 21% of service leaders reduced headcount due to AI.

- **[S11] Gartner (2026-11-04).** Survey: 4% realized expected AI ROI in customer service.

- **[S12] JD.com Investor Relations (2026-03-05).** JD.com fourth quarter and full-year 2025 results.

- **[S13] AIGW Observatory (2026-10).** JD.com: From Customer Service to AI Operations.

- **[S14] State Council General Office, PRC (2026-10-23).** Opinions on adapting to AI technology to build employment-friendly development (Guobanfa 2026 No. 29).

- **[S15] State Council of China English website / Xinhua (2026-10-23).** China issues guidelines to promote employment amid AI advances.

- **[S16] EUR-Lex (2026-07-27).** Regulation (EU) 2024/1689, consolidated text, Article 50.

- **[S17] Reuters (2025-12-16).** Business leaders agree AI is the future; they wish it worked right now.

- **[S18] TechCrunch (2025-06-04).** Klarna CEO says company will use humans to offer VIP customer service.