

ada

The Generative AI Toolkit for Customer Service Leaders

Templates and recommendations to help you make the most of your AI Agent.

Paving the way for AI-first customer service

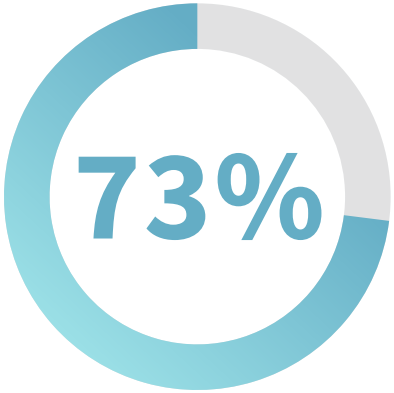
We predicted that in the near future, every company will be an AI company, and sure enough, we were right. Today, every company is working on incorporating AI into its core business processes.

For customer service, this means transitioning from an **agent-first** to an **AI-first** model, where the nucleus of the service organization is the AI Agent and leaders build around it. Your customers are expecting this too.

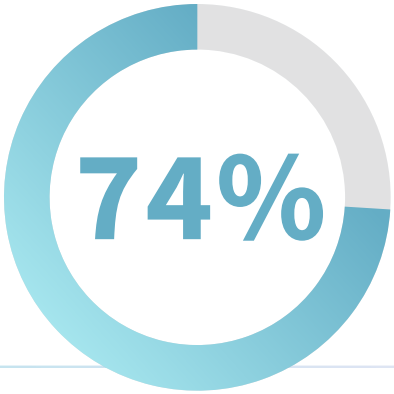
The most transformative effect of an AI-first model is elevating the role of the service organization as a whole.

While previously it may have been seen as a necessary cost center, Customer Service will now have the authority to influence core decisions that affect the growth of the company — decisions like how to evolve the product or service, which third-party vendors to partner with, what the brand should look and sound like…

Companies that embrace this shift will see the most growth.



of surveyed consumers expect more interactions with AI in their daily life



of surveyed consumers believe AI will improve customer service efficiency

What's inside?

1. Job descriptions and practitioner profiles
2. Transitioning to an AI-first Customer Service organization and example org charts
3. Preparing the customer service organization for an AI Agenr
4. Automated Resolution and other evolving customer service metrics
5. Key AI-first customer service metric benchmarks and how to analyze and report on them
6. Understanding the privacy, data, and security risks of an AI Agent
7. Legal and security questions to ask while assessing different AI Agent vendors

How to use this toolkit

These resources will help you understand and plan for the shift to an AI-first model. Of course, no two companies are alike, and you're welcome to adapt any of these resources to better suit your specific needs.

Starting from scratch? Use this toolkit as a foundation for onboarding an AI Agent and creating an early growth strategy. **Otherwise,** use it as a guide to optimize your efforts and long-term roadmap.



AI-first job descriptions and practitioner profiles

Recommended reading

- The people defining the CX careers of the future

AI Specialist

Job description

We're seeking a AI Specialist to join our Automated Customer Experience (ACX) team and leverage the capabilities of an AI Agent to deliver exceptional service to our customers.

In this role, you'll be responsible for onboarding our AI Agent and coaching it to increase our automated resolution across all our support channels. You'll also be in charge of surfacing gaps and opportunities in our knowledge base and support documentation, as well as customer insights that can guide the work of various organizations within the company.

No technical experience is necessary, instead we're looking for someone who has a deep empathy with our customers. The most successful candidates will have had experience working closely with different teams, on different product lines, or held different positions within the same company.

Qualifications

- 3-5 years of experience in customer service
- Strong writing and language skills
- Strong project management skills
- Strong communication skills

Responsibilities

- Work with our AI Agent partner to onboard the AI Agent
- Connect the AI Agent to software within our tech stack to enable it to take action on behalf of our customers
- Use pre-trained AI/ML algorithms to coach the AI Agent and increase its efficiency over time
- Optimize our knowledge base and other support documentation for the AI Agent
- Review and edit content generated by our AI Agent
- Review chat transcripts and customer service data to look for insights and opportunities
- Surface content gaps and opportunities for improvement in support documentation
- Use the AI Agent platform features to ensure that all generated answers adhere to our brand to create a cohesive experience for customers
- Report frequently on AI Agent performance and understand levers to pull to optimize against company goals

AI Manager

Job description

We're seeking a AI Manager to join our Automated Customer Experience (ACX) team and help us deliver exceptional service to our customers using generative AI.

In this role, you'll be responsible for onboarding our AI Agent and coaching it to increase our automated resolution across all our support channels. You'll also be leveraging the insights and analytics dashboard to optimize the automated customer service experience and developing strategies that can widen the scope of the AI Agent's abilities and value to the organization.

We're looking for creative problem-solvers with experience managing cross-functional projects. The most successful candidates are curious about how customers interact with different parts of the company, and can use customer insights to influence the roadmaps of various teams.

Qualifications

- 5-8 years of experience in customer service
- Strong data analysis skills
- Strong project management skills
- Strong communication and presentation skills as well as the ability to organize and implement requirements from multiple cross-functional teams
- Technical experience with API's preferred
- Experience with and/or deep curiosity for NLP, customer intent recognition, human-centered design, interaction design, information architecture, user research, and/or content design
- Experience with customer service automation is a bonus

Responsibilities

- Ensure the AI Agent is automatically resolving customer inquiries consistently and to the best of its ability
- Drive ongoing improvement projects for the AI Agent
- Collaborate with stakeholders to identify key integrations for a more personalized automated experience
- Review customer analytics and audit transcripts to identify optimization opportunities
- Utilize data, CS experience, and market insights to influence the company's content library and strategy
- Identify customer problems, uncover patterns, discover opportunities, and visualize possibilities to increase automated resolutions
- Propose solutions to address gaps and act on opportunities, and collaborate with partners and stakeholders to see them through
- Report frequently on the AI Agent's performance and understand levers to pull to optimize against company goals
- Work with Product Managers, Data Analysts, UX Designers, Back-end Engineers, and Customer Support Agents to create better experiences across all aspects of a customer's interaction with the company

Director of ACX

Job description

We're looking for a senior customer service leader to head up our Automated Customer Experience (ACX) team.

In this role, you'll be working with other leaders from teams like Product, Engineering, Data Science, Business Intelligence, Marketing, and more to align on company goals and design strategies that leverage the AI Agent to meet them. You'll also be managing our talented AI Specialists and AI Managers.

Your work will have a marked impact on the company and you'll have an opportunity move the needle by influencing business decisions.

Qualifications

- 8+ years of experience in customer service
- 3+ years of experience as a customer service leader
- Data science and analysis
- Strong management and project management skills
- Strong communication and presentation skills as well as the ability to organize and implement requirements from multiple cross-functional teams
- Experience designing short term and long term customer service strategies
- Experience collaborating with multiple cross functional teams on a common project
- Ability to align business objective to team goals and priorities
- Experience with customer service automation and familiarity with different software in the customer service tech stack

Responsibilities

- Develop short term and long term roadmaps for AI-first customer service
- Work with teams across the organization to integrate the AI-first customer service strategy into their own workflows to assess and prioritize opportunities and constraints — advocating for the end user while balancing priorities and business goals
- Define ACX team workflows, establish project timelines, implement process improvements, and draft improvement project proposals with key stakeholders
- Manage and measure a growing ACX team, including strategy for work stream alignment, hiring, onboarding, coaching, and on-going development
- Create quarterly OKRs that align with overall company goals, and report on the ACX team's value and impact
- Identify successes and areas of opportunity for the AI customer service experience, promoting the work of the ACX team
- Surface trends to help inform key changes and improvements across the entire company
- Maintain in-depth technical knowledge of automation platform(s) workflows, dependencies, and limitations
- Deliver customer insights to business intelligence teams

Practitioner profile - AI Specialist



Heather Bryant

System Admin

Day-to-day:

- Working with the AI Manager to find creative ways to leverage Ada's applications and tools to better serve the organization
- Reviewing upcoming releases to Intuit Mailchimp's products that may require content updates and assigns projects based on this
- General AI Agent coaching and maintenance

How to level up:

- Data science: how to uncover issues and opportunities from customer service data
- Project management: how to drive the implementation of projects related to the AI Agent
- Deeper understanding of APIs: how to make the most of the connections between all the different software in the customer service tech stack

Heather's advice

I find the most useful skill is being able to ask the right questions so I can get the right answers.

Practitioner profile - AI Manager



Edgar Lara

ACX Manager

Square

Day-to-day:

- Evaluating Ada's applications and tools to figure out how Square can leverage the platform to better serve the needs of cross-functional teams
- Creating thoughtful automations that anticipate customer needs and improve the overall operations of the business
- Analyzing customer data to uncover problems
- Proposing solutions to customer service issues and flagging areas of opportunity for Product teams

How to level up:

- Strategy: taking a customer-focused approach to problem solving while prioritizing different updates and optimizations
- Business storytelling: building proposals, reporting on success, and getting stakeholder buy-in

Edgar's advice

It may seem simple at first, but you're stepping into this emerging world of AI and conversational design, and things are still evolving. **You really have to be open to change, taking risks, and trying something new.** That's where you're going to have the most fun while still making a significant impact.

Practitioner profile - Director of ACX



Sarah Ricketts

Director of Business Solutions

INTUIT mailchimp

Day-to-day:

- Managing the ACX team
- Goal-setting, reporting success, and creating a long-term vision for the ACX team
- Designing strategies that use customer service automation to improve the company overall
- Aligning with senior leadership
- Reporting on the success of the team

How to level up:

- Deep understanding of how generative AI is changing the customer service landscape
- Pulling the right levers to set up the company for a transition into AI-first

Sarah's advice

Set up your team for success by giving them clear job responsibilities. A Product Owner *[i.e. AI Manager]* allows us to partner with our business owners to understand their needs and our product roadmap. A Systems Administrator *[i.e. AI Specialist]* focuses on making it happen.



Transitioning to an AI-first Customer Service organization

Keep in mind...

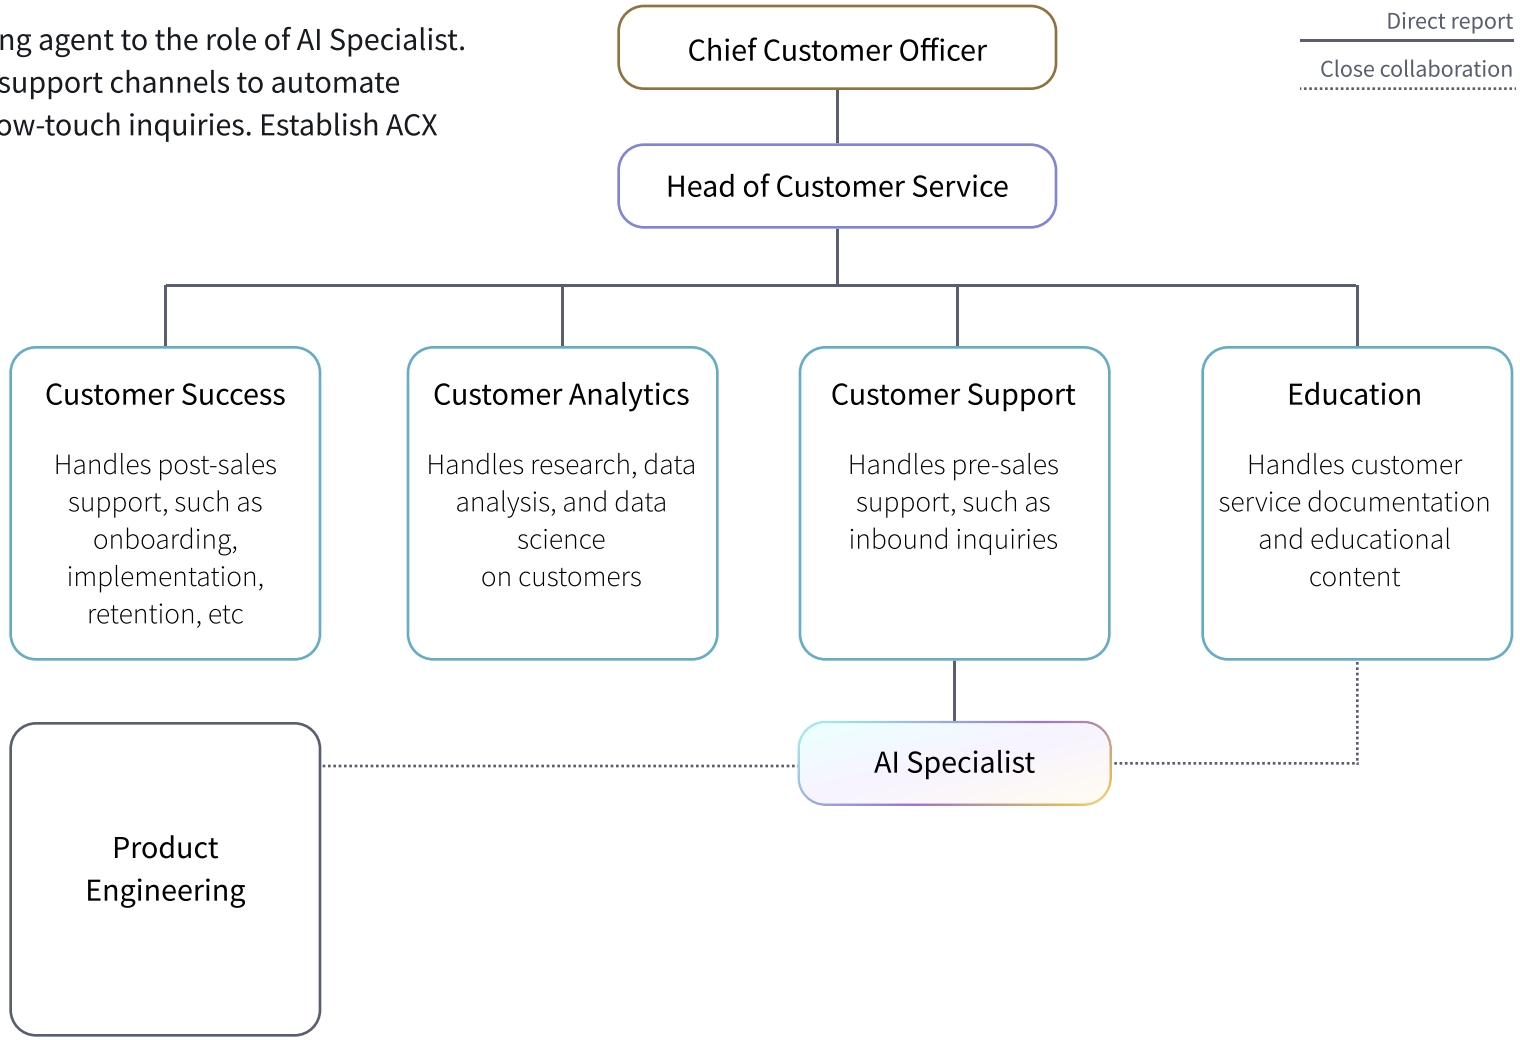
These org charts represent one example of a Customer Service org architecture and is meant to offer broad strokes for you to use as a starting point. We work with customers who have different structures that vary depending on company size and industry — centralized, decentralized, matrix, journey-led... Your own company may have a different structure as well.

As you scale through the phases, you'll be able to customize the org chart to better suit the company's needs. You may find that you need more AI Specialists or AI Managers, more hierarchy, or that you might have the ACX team report to a different leader.

There are no right or wrong answers, as long as the structure you're building empowers the team and sets them up for success in an AI-first world.

Phase 1: Onboard

Transition your most promising agent to the role of AI Specialist. Onboard the AI Agent across support channels to automate resolutions of high-volume, low-touch inquiries. Establish ACX benchmarks and KPIs.



Transitioning to AI-first

Ensuring success in phase 1

People

- Onboard a AI Specialist focused on AI Agent onboarding and content management, reporting to the Customer Support leader
- Introduce the AI Specialist to the wider Customer Service org and explain their role
- Introduce the AI Specialist to the Education team as a major stakeholder in Knowledge Base content management as well as internal documentation on AI Agent interactions (e.g. the relevant workflows on handoffs from AI Agent to support agent)
- Introduce the AI Specialist to the AI Agent platform partner as the person who will be onboarding, coaching, and managing the AI Agent, as well as learning from the partner's best practice recommendations and generative AI expertise
- Introduce AI Specialist to the Product/Engineering teams for integrations and embedding the AI Agent across the website

Technology

Connect the AI Agent to:

- Customer service agent platform
- Knowledge Base tool

Strategy

- Ensure the Knowledge Base documentation is accurate, up to date, and optimized for generative AI *(more on this in the next chapter)*
- AI Specialist meets weekly with the Education team to surface any content gaps or opportunities gleaned from reviewing the AI Agent transcripts
- Set quarterly goals for AI Agent as part of the Customer Service OKRs
- AI Specialist provides monthly reporting against the goals to their manager

Automated interactions

- Automated FAQ
- Handing off a conversation to a human agent when necessary

Business impact

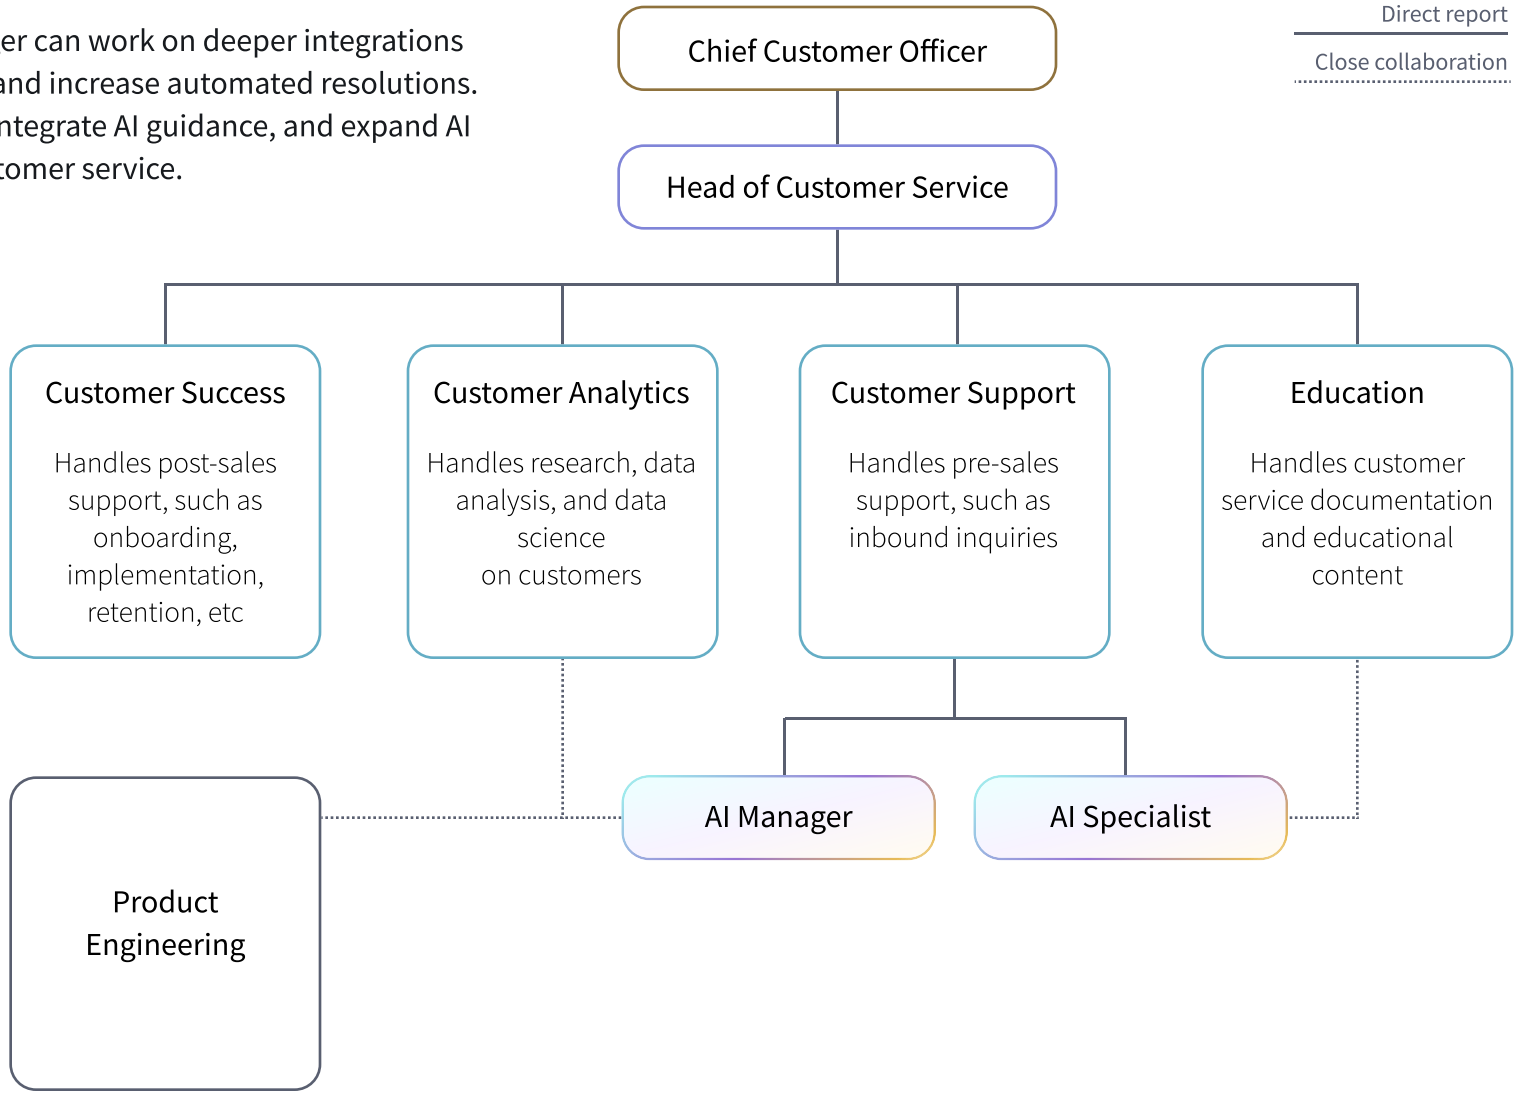
- Deflect from higher cost to lower cost channels (e.g. phone to chat)
- Significant agent capacity freed
- Improved agent productivity from assisted resolution mechanisms such as improved context during handoff and pre-authentication

Key metrics

- Lower agent headcount
- Lower cost per conversation
- Lower overall support costs

Phase 2: Measure and coach

A strategic, creative AI Manager can work on deeper integrations that power complex actions and increase automated resolutions. Leverage machine learning, integrate AI guidance, and expand AI Agent capabilities within customer service.



Transitioning to AI-first

Ensuring success in phase 2

People

- The AI Specialist continues to focus on day-to-day AI Agent coaching and content optimization
- Onboard an AI Manager focused on customer insights and automating more complex resolutions, reporting to the Customer Support leader
- Introduce the AI Manager to:
  - The Analytics team to analyze the customer data from the AI Agent dashboard and uncover insights
  - The Engineering team to integrate the AI Agent into more systems and power more complex actions, as well as potentially making UI changes, deploying on different channels, embedding on different pages, etc...
  - The Product team to report customer problems, trends, and propose solutions to improve
  - The AI Agent platform partner to continue education past phase 1 (e.g. CSM and/or Professional Services beyond implementation for continuous improvement.)
  - If relevant, other key third party stakeholders such as CRM platform partner(s) and/or contractor(s)

Technology

- Integrate the AI Agent to the CRM and connect it to any necessary software via API endpoints

Strategy

- Configure chat to be primary support channel (suppress phone number and email)
- Adjust quarterly goals for AI Agent as part of the Customer Service OKRs
- AI Manager provides monthly reporting against the goals to their manager
- Publicize AI generated insights across the company to get other teams thinking about how they can leverage the data
- Publicize AI Agent capabilities across the company to get other teams thinking about how they can also be AI-first

Phase 1 plus...

Automated interactions

- Customer authentication
- Personalization
- Self-service (e.g. order tracking, password reset)
- Actions (e.g. refund request, book flights)

Business impact

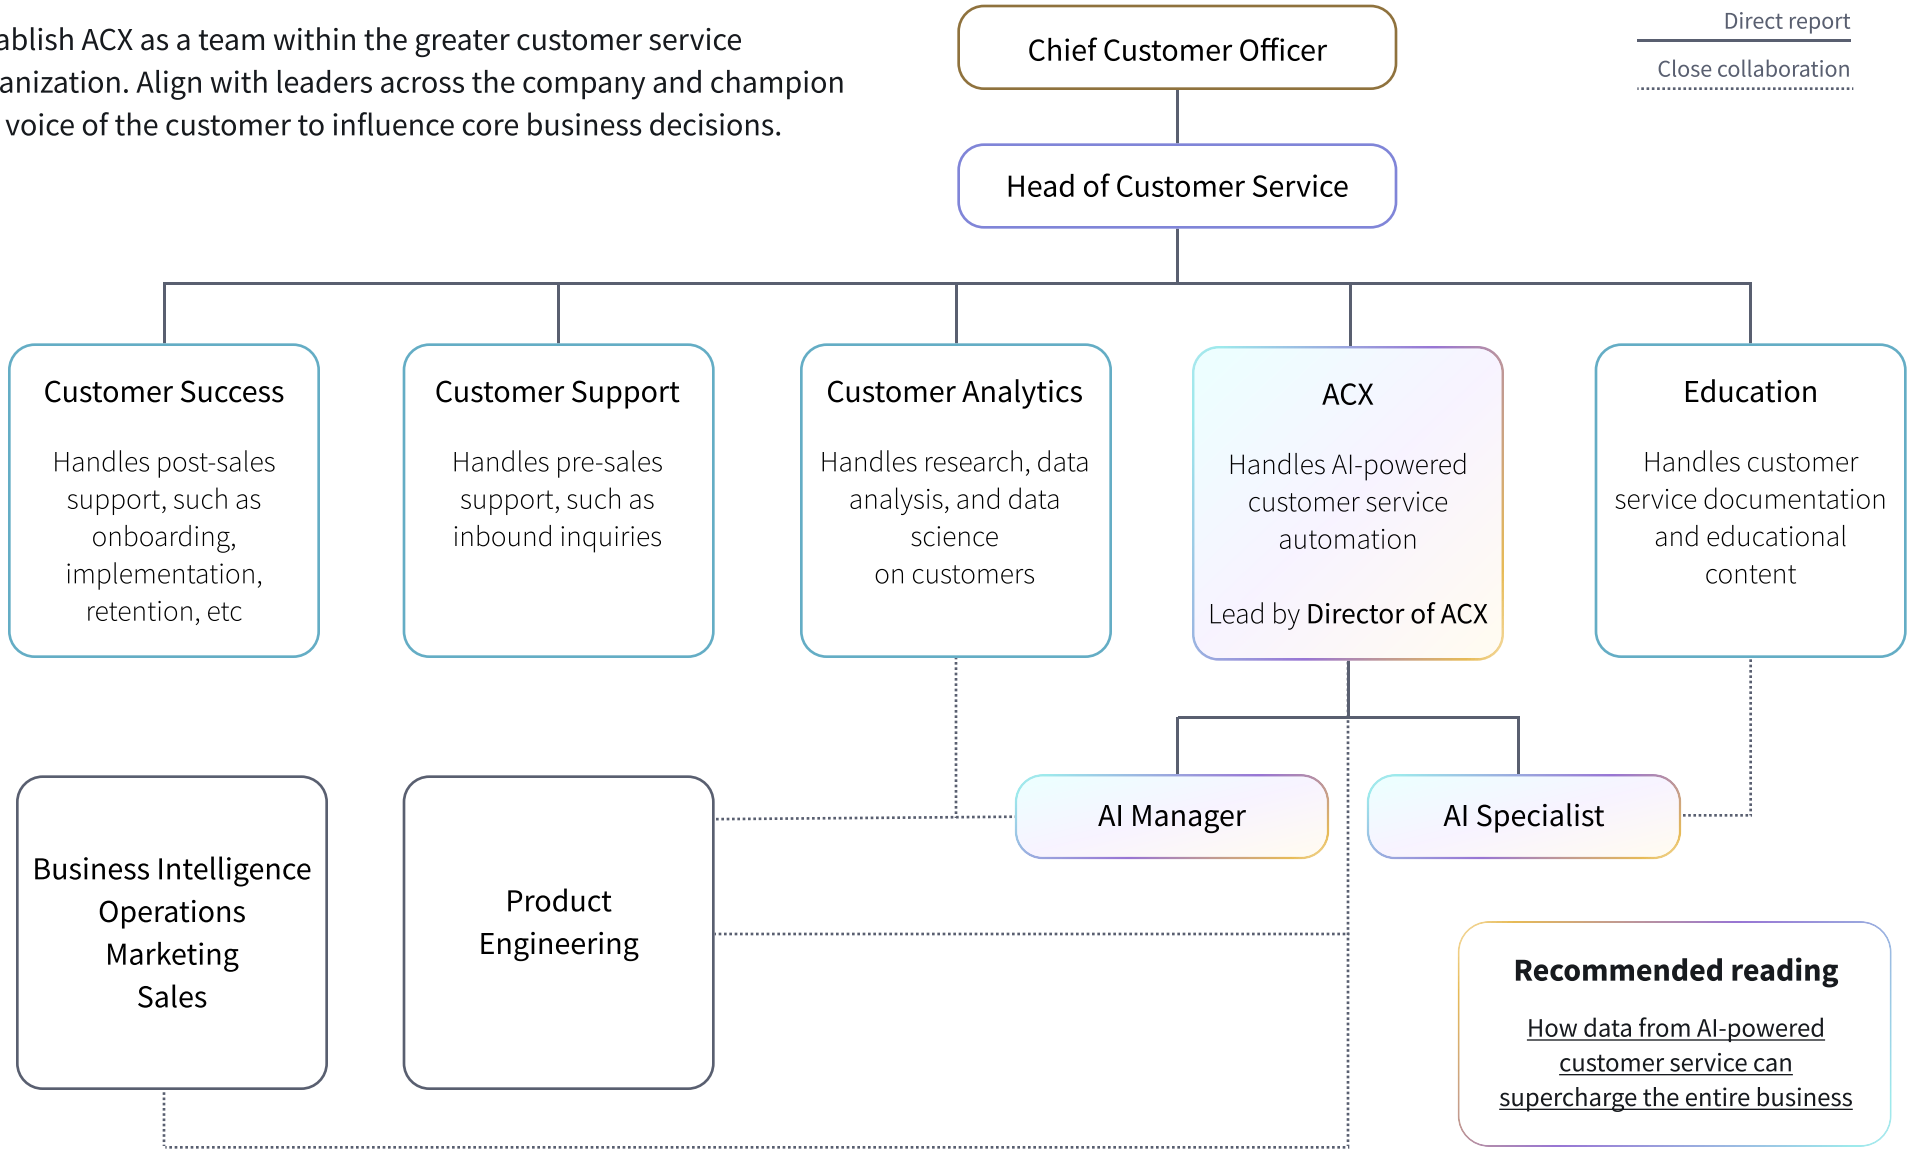
- Customer service becomes a competitive differentiator through improved response/resolution time and personalization
- Happier customers lead to repeat purchases and expansion revenue
- AI Agent makes expanding to different geographies possible without the need to hire more agents

Key metrics

- Increased automated resolutions
- Lower avg response & resolution time
- Higher CSAT/NPS
- Lower churn & higher expansion revenue
- Reduced time to expand support to new geographies

Phase 3: Promote

Establish ACX as a team within the greater customer service organization. Align with leaders across the company and champion the voice of the customer to influence core business decisions.



Transitioning to AI-first

Ensuring success in phase 3 and beyond

People

- The AI Specialist continues to focus on day-to-day AI Agent coaching and content optimization
- The AI Manager continues to focus on data and integrations
- Create an Automated Customer Experience (ACX) team
  - Onboard a Director of ACX focused on growing the influence of customer service on other teams and core business decisions
  - The AI Specialist and AI Manager move to the new team and report to the Director of ACX
- Introduce the Director of ACX to:
  - Product and Engineering leaders to design a cohesive API strategy that powers the AI Agent to take the same actions as a human agent
  - Business Intelligence and Customer Analytics leaders to analyze and report on insights and trends gleaned from AI Agent data
  - Sales, Marketing, and Ops leaders to design strategies that use the AI Agent and its data to achieve company goals
  - Other key leaders such as those at partner companies (e.g. CRM, other systems in tech stack)

Technology

- Integrate the AI Agent to all internal systems and endpoints

Strategy

- Adopt an omnichannel approach to customer service automation
- Adjust quarterly goals for AI Agent as part of the Customer Service OKRs
- Director of ACX reports on the ROI of AI-powered automation
- Uncover new ways to take an AI-first approach in other parts of the company

Phase 2 plus...

Automated interactions

- Ability to serve customers at the same level as human agent

Business impact

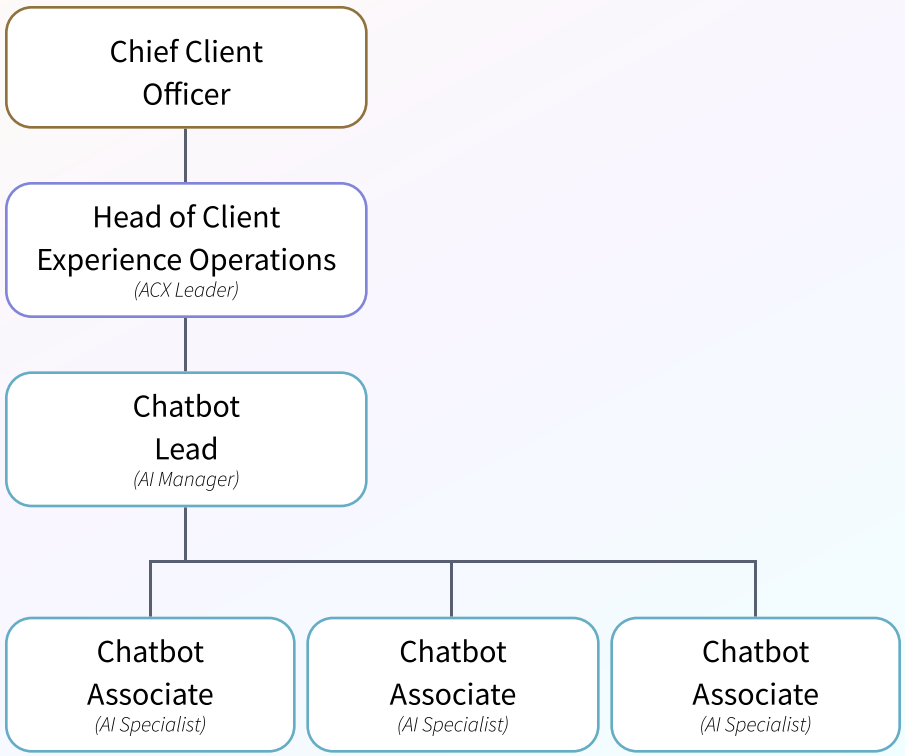
- Reduce or eliminate costs and friction from requiring technical resources to automate workflows and resolve inquiries
- Improved understanding of leading indicators for product/service/market performance
- Agent capacity can be redeployed towards revenue

Key metrics

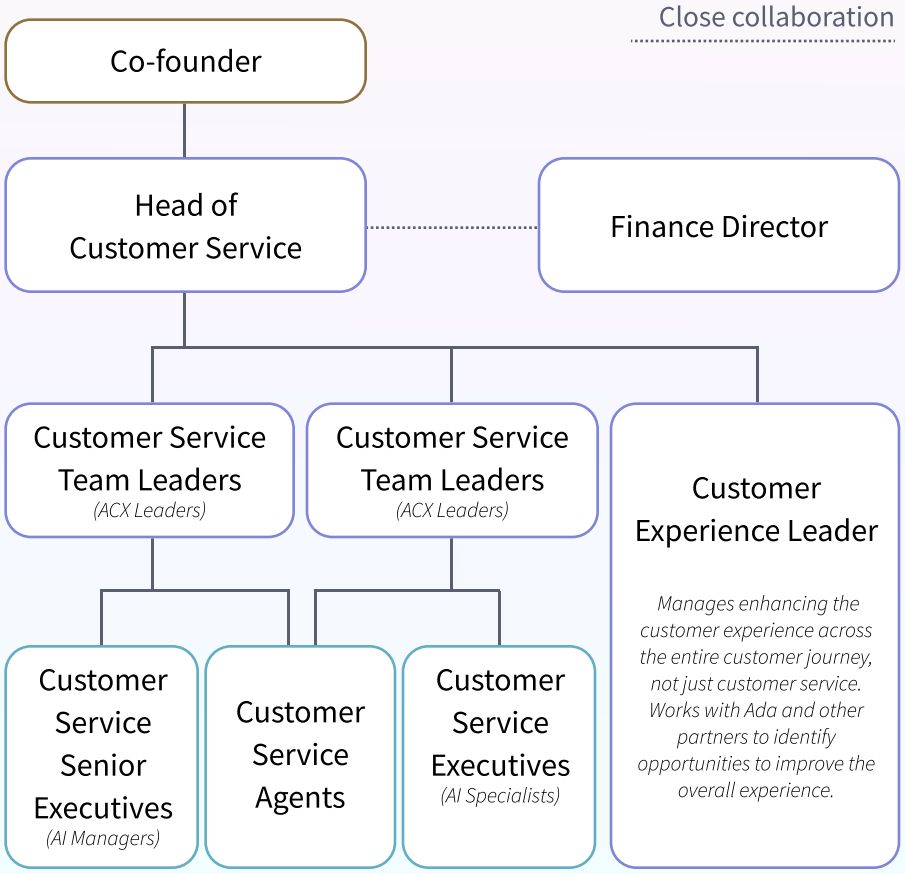
- Highest automated resolution potential
- Even lower agent headcount
- Higher agent satisfaction and employee retention
- Increase in expansion revenue
- Lower costs leads to higher budget to re-invest in customer service

Examples of ACX teams at companies using Ada

Wealthsimple



oh POLLY



Transitioning to AI-first

Future state

To truly harness the potential benefits of becoming AI-first, every single team on the Customer Service organization has to be empowered to use AI in their day-to-day to reach their goals.

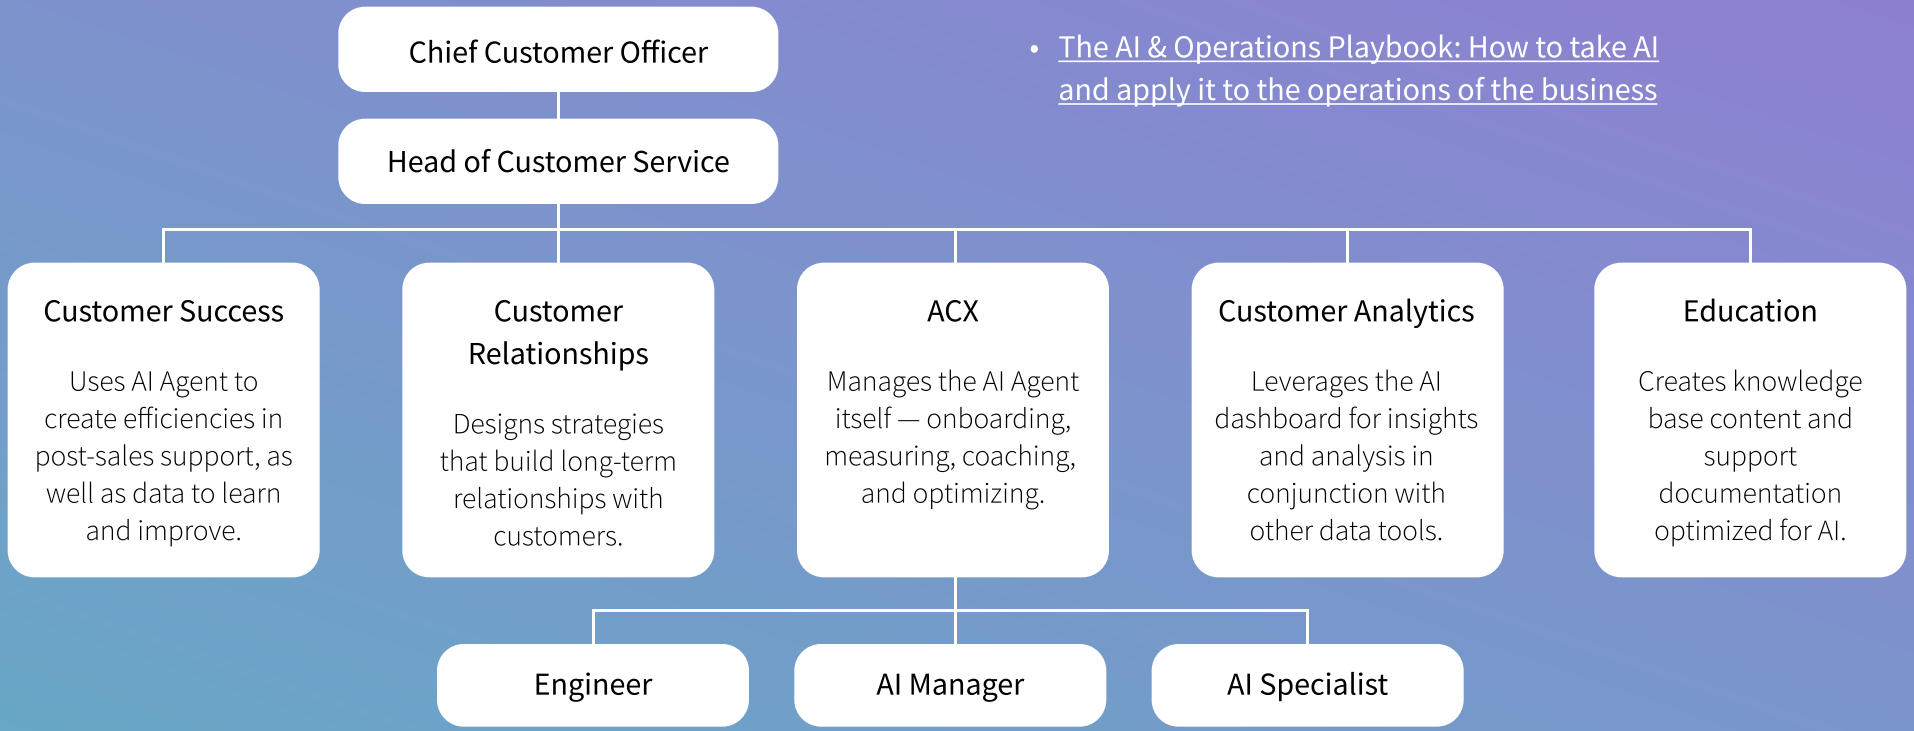
In the near future, AI Agents will be advanced enough to completely cover for human support agents, regardless of the team size, geographic location, or number of support channels. And while the future is agentless, it's not humanless.

The Customer Support team will transform into a Customer Relationships team, charged with designing strategies that take into account the entire customer experience, across all touchpoints — marketing, product, service, pricing, and more.

The CCO's role is to ladder up the goldmine of insights and proposals to the executive team to influence decisions across the entire company.

Recommended reading

- The AI & Operations Playbook: How to take AI and apply it to the operations of the business





How to prepare for onboarding an AI Agent

Before you start

This worksheet is meant to be used in conjunction with an on-demand webinar we hosted featuring a guest speaker from Forrester titled “How to Leverage Generative AI for Customer Support Automation”.

In order to make the most out of this worksheet, please watch the webinar first.

Watch the webinar →

Knowledge: Teach your AI Agent how to read

The AI Agent is going to be using support documents to generate information. These documents can be anything from knowledge base articles, product manuals, internal guides, or even historical chat transcripts.

Audit your documents and answer these questions to surface opportunities for improvement.

- Where does your support documentation live?
  - Is it easily accessible?
  - Is there a single source of truth?
- How often does your support documentation get updated?
  - How does it get updated?
  - Does it get updated globally across the entire organization or only local versions?
- Is your support documentation helpful?
  - Is it clear?
  - Is it thorough?
  - Is it up to date?
  - Is it accurate and truthful?
- Is your knowledge base optimized for AI?

Recommended reading

- Knowledge base best practices for generative AI

Action: Teach your AI Agent how to do things

In order for the AI Agent to automatically resolve customer problems, it needs to be connected with other systems from across your organization to access information such as customer or order data, and trigger actions such as preparing returns or resetting passwords. It's essential that you have your API strategy locked in.

Audit your customer service tech stack, then answer these questions.

- What software is currently in your tech stack?
  - How much of it is critical to customer service operations?
  - Is there any software that can be replaced with automation?
- Does all your critical software have an API?
  - What can the API do?
  - How well is it built?
  - How well is it supported by the vendor?
  - Is the API stable enough to support an application build on top of it?
- Knowing all this, is there software that needs to be updated or replaced?
- Are there software gaps you need to fill?

A few tech stack strategy recommendations

Prioritize software that your customer service automation platform:

- Integrates with out-of-the-box
- Has open endpoints
- Can solve for multiple use cases (e.g. Salesforce to push account level information and for case creation)
- Is accessible and used across multiple teams or departments
- Is easily customizable and configurable without developer resources

Example of an ideal SaaS tech stack

- Phone Support/IVR : Talkdesk
- Internal collaboration tool: Slack
- Channels: Web, App (iOS, Android), SMS (Twilio)
- Social Channels: Twitter, Facebook, WhatsApp, IG
- CRM: Salesforce
- Analytics Platform: Tableau
- KB Platform: Salesforce Knowledge Base
- Marketing Automation Tool: Pardot
- Ticketing: Salesforce
- Live Chat: Salesforce Live Agent

Example of an ideal Ecommerce tech stack

- OMS: Shopify
- Ecommerce / Payment system: Stripe
- Phone Support/IVR : Zendesk Talk
- Channels: Web, App (iOS, Android), SMS (Twilio)
- Social Channels: Twitter, Facebook, WhatsApp, IG
- Analytics Platform: Power BI
- KB Platform: Contentful
- Marketing Automation Tool: Klayvio
- Ticketing: Zendesk Ticketing
- Live Chat: Zendesk Live Chat

Expertise: Teach your AI Agent how to solve hard problems

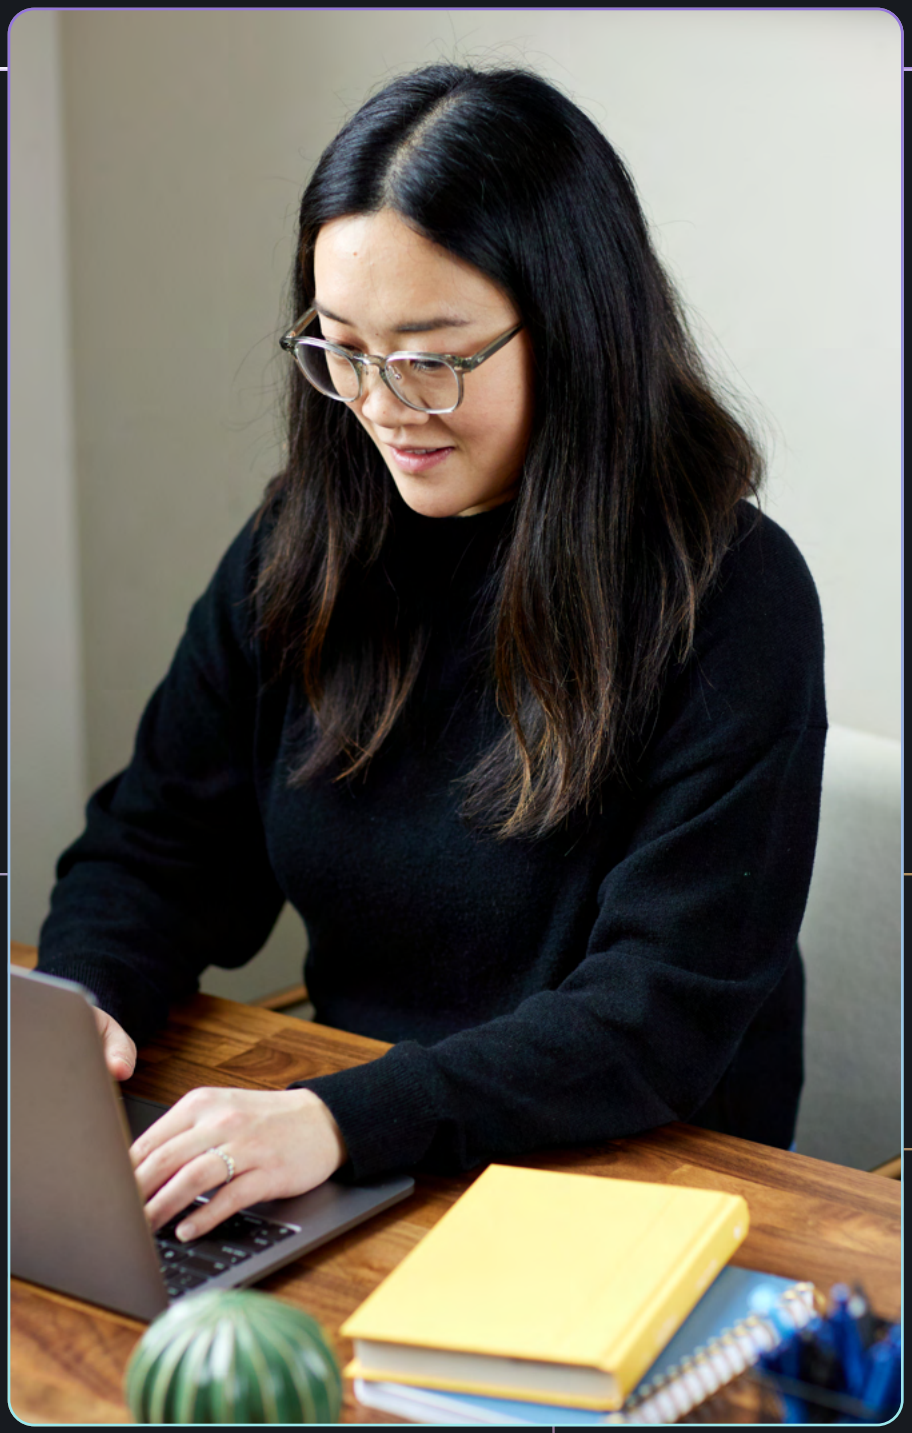
Humans are an important part of your AI strategy. Tapping into your human agents' expertise and making it a part of the AI Agent initiative will contribute to its success.

Map out your customer service organization and answer these questions to surface opportunities for improvement.

- Where do your sources of expertise sit within the organization?
- What is the expertise that they have?
- Is there expertise that's currently missing that you need to hire for?
- Are some of these experts currently outside the customer service organization?
  - If yes, how can you involve them?
- Are there star agents that you can promote to AI Manager?
- Consider who might need training and access to:
  - Review AI Agent interactions
  - Coach the AI Agent
  - Learn from insights
  - Surface business opportunities to other teams

Recommended reading

- Best practices for deploying generative AI in customer service



How to measure an AI Agent with Automated Resolution

Recommended reading

- The complete guide to customer service metrics for 2024

Automated Resolution

Evolving beyond agent-first metrics

In a few years, AI and automation will supplant human agents almost entirely, and human customer service agents will instead be elevated to a more strategic and important role in the company.

This foundational shift is naturally accompanied by a shift in how we measure success. Agent-first metrics will continue to be relevant — albeit adapted to a new paradigm — but there's a growing need to reliably measure the efficacy of AI Agents.

Historically, metrics like containment rate have come close to measuring a chatbot's success, but they fall short of providing a complete picture when it comes to AI Agents. The containment rate tells you which proportion of conversations that customers had with your bot ended without being handed off to a human agent. But without context, this metric can only tell you so much. Can you differentiate between a customer who ended a conversation with your bot because they were satisfied, as opposed to one who got frustrated and gave up?

It's easy to understand why current metrics make it hard to calculate the ROI of AI Agents, let alone the impact of the service organization as a whole.

So, let's take a step back and think, *"What's the goal/purpose of an AI Agent?"* The answer at the very core is *to automatically resolve customer inquiries*. This seemingly basic objective has a massive ripple effect.

AI Agents in customer service are only as good as the number of inquiries they can automatically resolve. It's as simple as that. The north star measurement should be Automated Resolution (AR).

Understanding AR

Automated Resolution (or AR) refers to a fully automated conversation between a customer and a company that is relevant, safe, and accurate and does not escalate to a human agent.

- **Relevant:** the AI Agent effectively understands the inquiry and provides directly related information or assistance
- **Accurate:** the AI Agent provides correct, up to date information
- **Safe:** the AI Agent interacts with customers in a respectful manner and avoids engaging in topics that cause danger and harm
- **No escalation:** the AI Agent resolves the inquiry without any agent support

The AI Agent assesses the conversation to determine whether a successful automated resolution did or didn't happen.

Measuring automated resolutions will revolutionize how you quantify your AI Agent's performance, helping you focus on what's most important to your customers.

Recommended reading

- How we optimize our AI for automated resolutions

Automated Resolution

How it works

Every hour, Ada takes a sample of engaged conversations that occurred 24 hours ago. A large language model assesses each conversation transcript for accuracy, relevancy, safety, and human involvement, then labels them as Resolved and Unresolved.

Resolved conversations indicate that:

- The conversation was accurate, relevant, and safe
- No human was involved

If your current AI chatbot solution does not do this automatically, you can still measure AR manually following this method:

1. Take a statistically significant sample of your chatbot's conversations
2. Review the transcripts and mark each conversation as **resolved** or **unresolved**, and **contained** or **uncontained**
3. Divide the number of **resolved and contained** conversations by the total sample number to get the Automated Resolution Rate (AR%)
4. Use this formula to calculate AR across your entire conversation volume:

AR = conversation_volume x AR%

Measure ROI using AR

AR makes it very easy to get a quick idea of the ROI on your AI Agent. Here are formulas for common customer service KPIs:

Agent hours saved
AR x average_handle_time / 60

Costs saved (agent hours)
hours_saved x agent_cost_per_hour

Cost savings (conversations)
AR x conversation_volume x cost_per_conversation

Projected headcount saved
AR / average_tickets_per_agent

Automated Resolution

Containment vs. AR

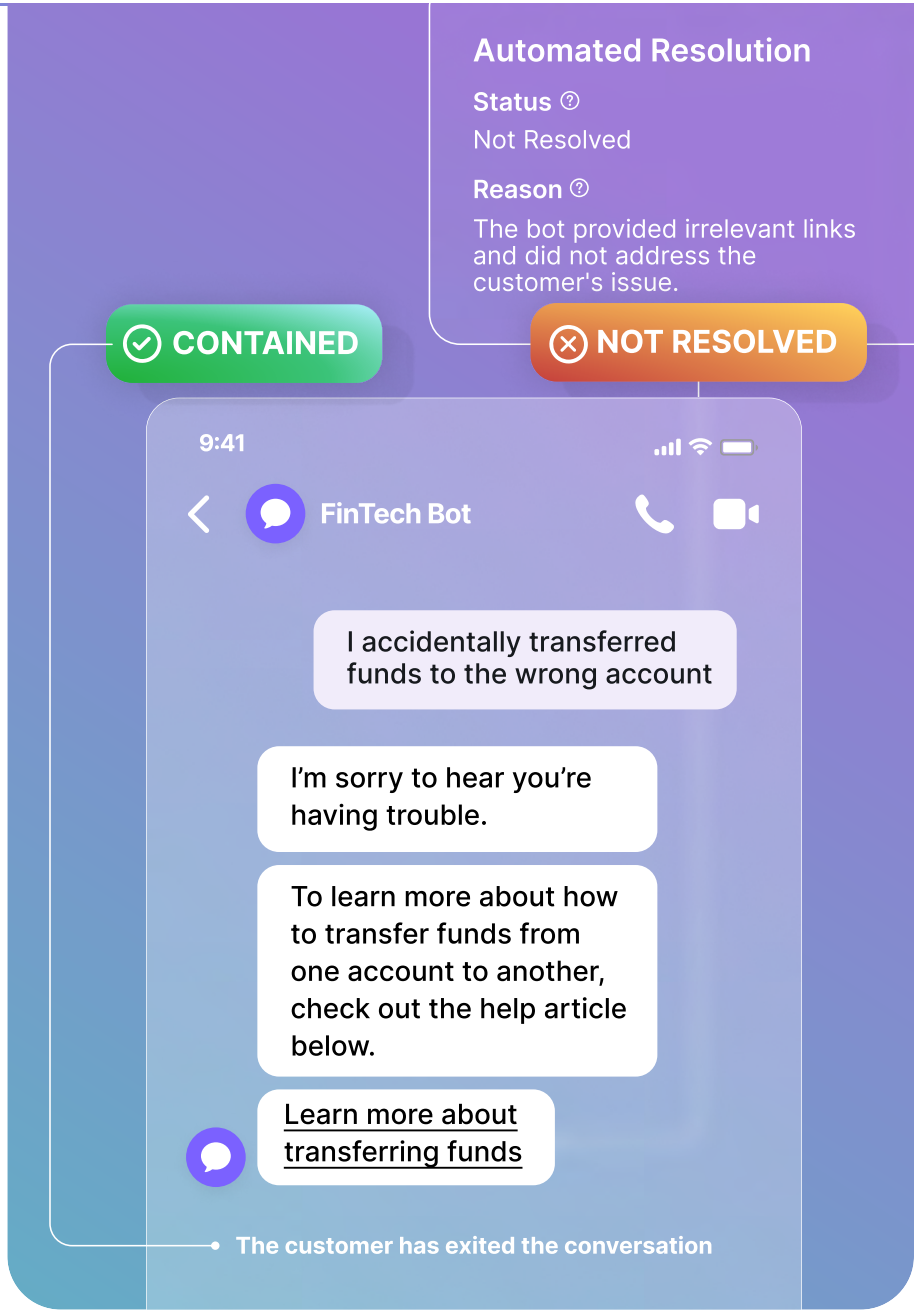
When you first start measuring AR, you'll naturally notice that the number is lower than your containment rate. This is normal.

Remember that containment tells you how many conversations didn't reach a human, regardless of whether or not the inquiry was resolved.

On the other hand, AR only takes into account **successful** conversations. AR is a more accurate indication of how helpful your AI Agent really is by measuring the success of those conversations and the content they contain. It gives you more insight, with the help of machine learning, to tell you if the customer got the help they needed.

AR also gives you insight into why a conversation was resolved or not, so you can more easily identify gaps and improve performance. With containment, it's almost impossible to identify why a customer exited the conversation (were they happy or frustrated?) so you think you're performing well, but you actually have problems you can't see and and can't fix their root cause.

The industry-standard metric of containment is only one piece of the puzzle. Automated resolution rate is more holistic, and is a better indicator of your bot's performance.



AI-first customer service metrics

From agent-centric to customer-centric

Currently, most of the KPIs that CS leaders track are agent-first and rely quite heavily on the performance of teams of human agents. The thing is, these KPIs are mostly influenced by one lever: budget. If you can hire more agents, have around-the-clock shifts, offer multi-lingual support, of course stats like CSAT, customer wait times, employee engagement, and more, will improve. But as we all know, budget is always a limiting factor.

AI-first customer service takes a more customer-centric approach. As more companies make the shift, CS leaders will need to introduce new KPIs for their teams that focus less on measuring human agents and more on measuring the quality of the service as a whole.

That doesn't mean that the current KPIs will disappear altogether, but they will evolve to suit the new way of working. Here are some examples.

AR will become the North Star

Increasing AR and AR% will become one of the biggest objectives of the CS org. CS leaders will design strategies for their teams to improve AR and increase conversation volumes.

AI Managers and AI Specialists will have new KPIs that directly influence AR.

Human agent metrics will become AI Agent metrics

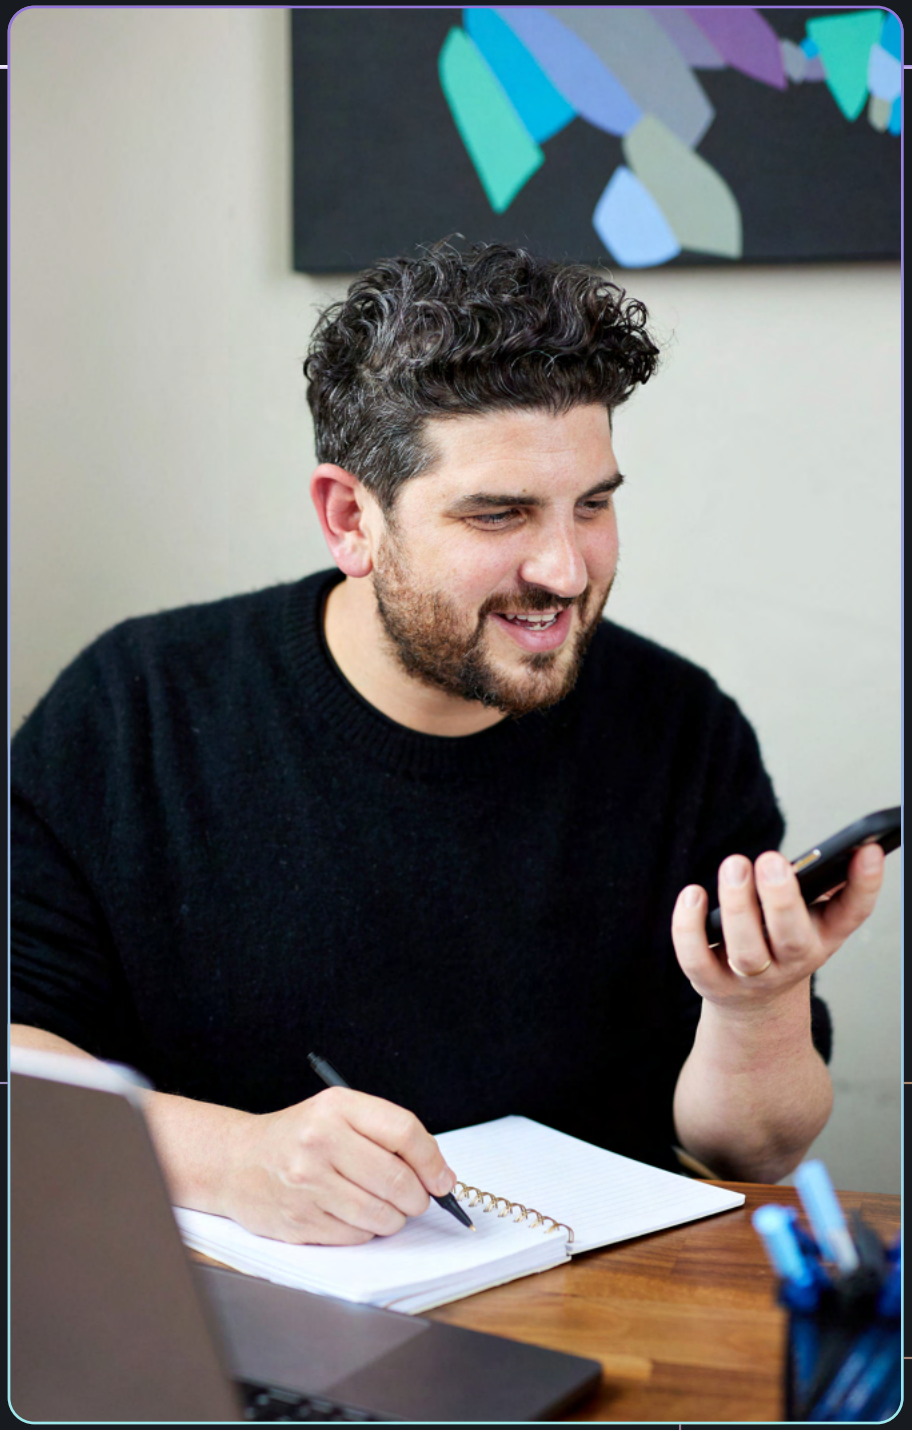
When companies reach 100% AR (meaning the AI Agent is resolving 100% of customer inquiries):

- CSAT becomes exclusively an AI Agent metric
- Average Handle Time will matter less from a cost perspective and more from an optimization perspective
- Cost per Resolution will more-or-less stabilize and will not factor in agent headcount
- First Contact Resolution will be a good indication of whether the AI Agent is able to accurately understand and resolve a customer's inquiry from the first try or if it needed more clarification
- CS leaders will likely use Time in Bot Dashboard as an indicator of AI Manager and AI Specialist engagement

Customer Experience agents will have revenue KPIs

The Customer Experience team (previously support agents) has a mandate to focus on long-term relationships with customers. To measure that, they'll have KPIs associated with revenue, such as:

- Return customers
- Customer loyalty
- Expansion
- Net Promoter Score



Key AI-first customer service metrics: analysis, reporting, and industry benchmarks

Keep in mind...

These benchmarks offer a quick glance at how your metrics compare against industry averages. They're intended as guidance, and may vary based on the size and maturity of a company.

*Benchmarks compiled in Q2 2024

Customer service data benchmarks

Automated Resolution Rate

Definition: the percentage of customer conversations that have been fully automated, that were relevant, safe, and accurate and not escalated to a human agent.

Some ways to learn from AR%:

- This is the best metric to determine how well your AI Agent is doing.
- Use this number to calculate the ROI of AI-powered customer service automation.
- Use this number while calculating headcount projections — either seasonal, full time, or expansion to new markets — by determining the volume of conversations that agents would have to handle.
- Dig into transcripts of unresolved conversations to see how you can optimize the AI Agent.
- During conversation volume spikes, check if AR% remained stable. If not, then you might need to plan ahead to create knowledge base content for seasonal spikes, major campaigns, product launches, etc...

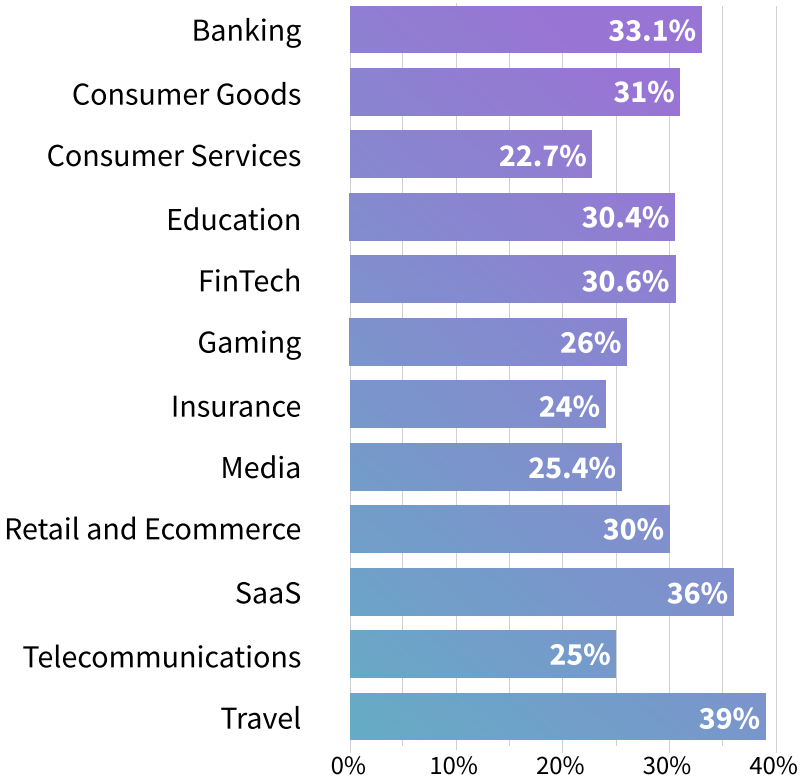
Check how AR% might influence these other metrics:

- CSAT Score: CSAT Score should increase as AR% increases.
- Average customer wait times for live agent: since the AI Agent is resolving more inquiries, wait times should decrease.

Average Automated Resolution Rate across all industries

32.5%

Average Automated Resolution Rate by vertical



Customer service data benchmarks

Containment Rate

Definition: the percentage of customer conversations that were not escalated to a human agent, regardless of whether or not they were resolved.

Some ways to learn from Containment Rate:

- CR% gives you a very quick glance at how many conversations are not being escalated to a human agent so you can quickly spot irregularities between specific time periods (day-to-day, week-over-week, etc...)
- Dig into conversations that were “contained” but not “resolved” to uncover opportunities for training and improving the AI Agent.

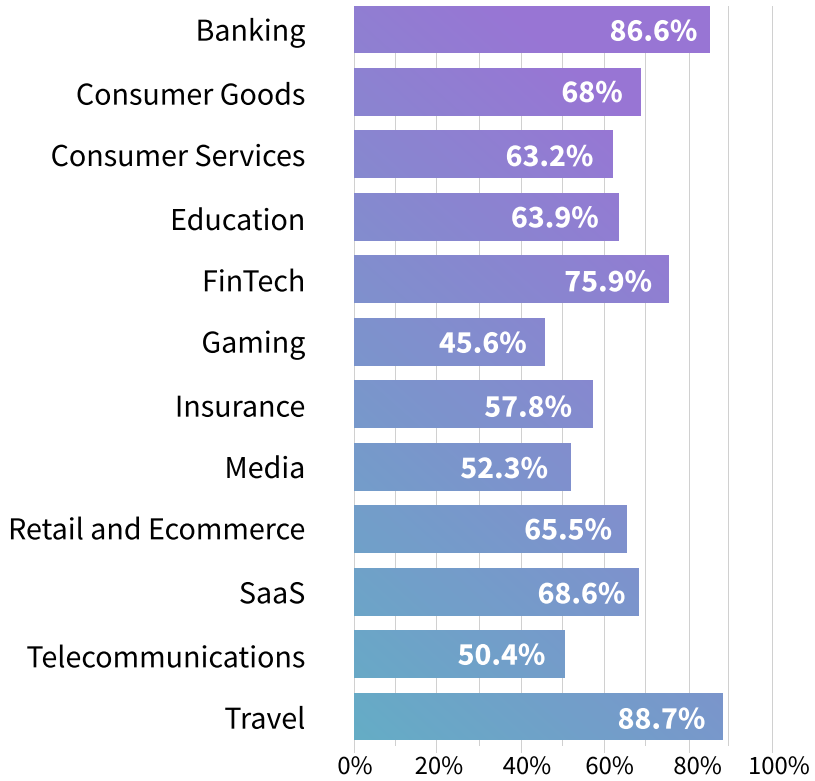
Check how CR% might influence these other metrics:

- AR%: keep an eye out on the discrepancy between AR% and CR%. While AR% will naturally be lower at the start, it should steadily increase and the discrepancy should decrease.
- Conversation Volume: as CV increases, CR% should remain stable. If it’s not, you can find opportunities to optimize automation by looking for patters in uncontained conversations.

Average Containment Rate across all industries

63.9%

Average Containment Rate by vertical



Customer service data benchmarks

CSAT Score

Definition: the percent of conversations customers reviewed positively, out of all conversations they reviewed.

Some ways to learn from CSAT Score Rate:

- Dig into the transcripts of conversations marked positively or negatively to uncover patterns that made the customers happy or unhappy.
- Be mindful of this metric specifically during big campaigns or seasonal sales, it can be a good indicator of whether the big brand moments were received positively or negatively by customers.

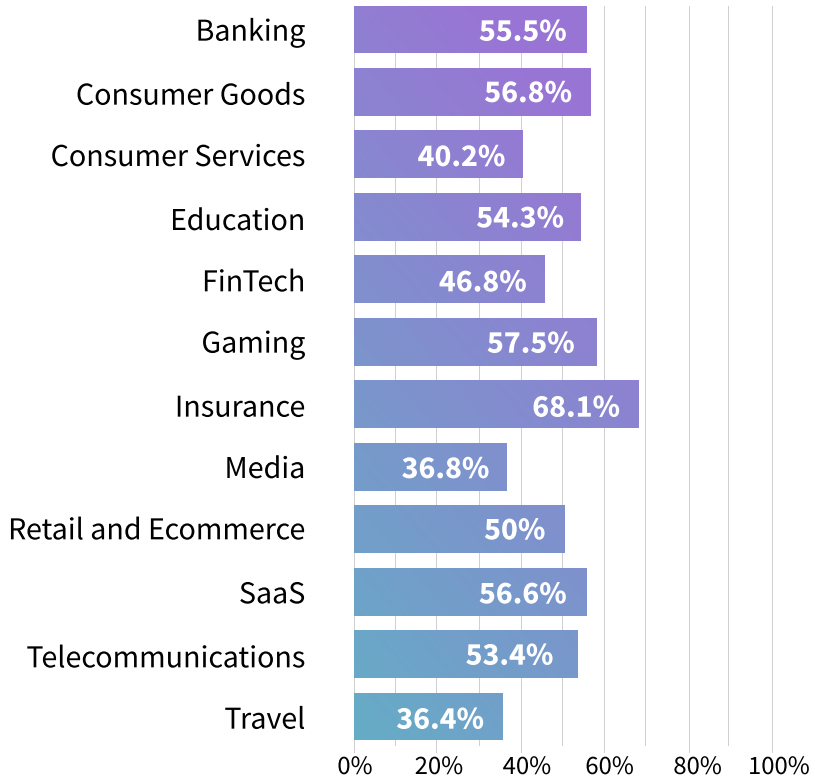
Check how CSAT Score might influence these metrics:

- AHT: CSAT Score typically increases as AHT decreases.
- AR%: CSAT Score typically increases as AR% increases.
- CR%: Comparing CSAT Score with CR% can indicate whether or not customers are happier when their inquiry gets escalated to a human agent.

CSAT Score across all industries

51.6%

Average CSAT Score by vertical



Customer service data benchmarks

Average Handle Time

Definition: the amount of time it takes to resolve an inquiry. This is measured separately for AI Agents and human agents.

Some ways to learn from AHT:

- For AI Agent AHT:
  - The AHT of the AI Agent should be proportional to the complexity of the resolution it's trying to automate — dig into the longest and shortest conversations to see if you notice patterns.
  - If you notice that AI Agent AHT for simple FAQ interactions is high, then you probably need to optimize your knowledge base content.
- For human agent AHT:
  - Review transcripts for conversations that are way over the average human agent AHT to see if agents need to be better enabled or if you can use the AI Agent to collect key information from customers before they get passed on.
  - Review the conversations with the shortest Agent AHT to see if they are better suited for the AI Agent.

Check how AHT might influence these metrics:

- CSAT Score: CSAT Score typically increases as AHT decreases.

Average Handle Times across all industries (min:sec)

AI AGENT:	HUMAN AGENT:
01:58	14:00

Average Handle Time by vertical (min:sec)

	AI	HUMAN
Banking	03:34	14:25
Consumer Goods	02:06	15:42
Consumer Services	02:15	15:07
Education	01:55	15:42
FinTech	02:25	15:31
Gaming	02:05	15:49
Insurance	01:28	12:14
Media	02:05	15:49
Retail and Ecommerce	02:24	13:14
SaaS	01:45	14:45
Telecommunications	02:03	14:58
Travel	02:25	15:12

Customer service data benchmarks

Conversation Volume

Definition: the number of conversations that customers have with a company per set length of time (day, week, month, quarter, year).

Some ways to learn from CV:

- Review yearly CV numbers to uncover patterns for volume spikes and be better prepared for them.
- Keep an eye on CV during big brand moments such as product launches or major campaigns to measure customer engagement.
- Unusual spikes in CV can signal outages or issues in other areas of the company, such as logistics or operations, and you can often notify the relevant teams before they are even aware of the issue.

Check how CV might influence these metrics:

- CSAT Score: comparing CV and CSAT Score during volume spikes can signal positive or negative brand sentiment.
- AR%: AR% should either increase or remain stable as CV increases. If AR% decreases, it means there are opportunities to improve the automated experience.

Average Conversation Volume across all industries

864 per day

Average Conversation Volume per day by vertical

Banking	2,084
Consumer Goods	522
Consumer Services	555
Education	236
FinTech	497
Gaming	1,883
Insurance	129
Media	1,192
Retail and Ecommerce	673
SaaS	564
Telecommunications	572
Travel	4,630

Example weekly AI Agent report

Weekly AI Agent metrics

Be sure to include week-over-week charts for these metrics to track trends

AR	AR%	CR%	AHT	CSAT Score
How many conversations did the AI Agent automatically resolve? XX	*What is the latest AR%?* XX%	*What is the latest CR%?* XX%	*What is the latest AHT?* XX	*Surveys collected* XX *CSAT Score* XX%

Weekly AI Agent insights

*Top 3 customer inquiries that were **automatically resolved***

Inquiry volume	AR%	Reasons for success

Weekly AI Agent insights

*Top 3 customer inquiries that were **escalated to a human agent***

Inquiry volume	Hand-off %	Opportunities for improvement

Other noteworthy insights

For example, any issues with the AI Agent or any knowledge base gaps or opportunities that the AI Agent identified.

ada

Example weekly AI Agent report

Updates and amendments

Such as:

- *Any improvements made to the knowledge base, product, or any other areas of the business that was informed by AI Agent data*
- *Any improvements made to the AI Agent itself — new integrations, new training, etc...*
- *Any issues from the previous weeks that have been resolved*

Weekly AI Agent successes

Such as positive customer reviews, early issue identification, anecdotes of very successful generations, or the AI Agent handling an unexpected CV spike.



Understanding the privacy, data, and security risks of AI Agents

Keep in mind...

The considerations outlined in this section are the most notable ones. It's important to note that these lists are not exhaustive, and are constantly evolving.

Recommended reading

- Is customer data safe in a world run by AI agents?

Definitions

Key terms that come up with regards to privacy, data, and security

Generative AI is a broad type of AI that focuses on creating new content, such as text, images, or music, by learning patterns and structures from existing data.

LLMs (Large Language Models) are mathematical model of human language that are often used within generative AI systems and power AI Agents, but that also have other non-generative applications such as sentiment analysis, text classification, authorship attribution and authenticity.

Input refers to the query that the customer asks the AI Agent.

Output refers to the response generated by the LLM.

Generative reply or generated output refers to the reply that is returned to the customer.

AI Agent refers to AI-powered customer service automation platform that you would use to automate customer service (e.g. Ada, the product).

AI Agent provider (and/or vendor) refers to the company whose platform you're using (e.g. Ada, the company).

LLM provider refers to the company that develops and offers the LLM for other third parties to utilize (e.g. OpenAI).

Customer refers to the end-users chatting with the AI Agent.

Recommended reading

- Chatbot vs. AI Agent: What's the difference and why does it matter?

Understanding the input risks

Risks associated with the data that is being provided to the AI Agent.

Risk	Definition	Risk to your customer	Risk to your company
Sensitive data exposure	Any input that customers provide to an AI Agent could be exposed to unauthorized entities, or the public internet. Even with contractual agreements in place that are meant to protect sensitive data, there is no way to fully account for human error and data leaks, intentional or otherwise.	Without proper safeguards, the customer's data may be used by your company, the AI Agent provider, or the underlying LLM provider for purposes other than just the AI Agent.	Legal, financial, and reputational repercussions resulting from the leak or misuse.
Privacy and profiling	Any input that customers provide to an AI Agent could be used to infer data about, or build a profile around, the person providing the input. Inferred data or profiling may include things such as personal preferences, beliefs, or demographics. The aggregate input data can also be used to build a profile about your company, including details such as web traffic data, order volumes, customer demographics, or business strategy.	Without proper safeguards, the customer's profile may be used by your company, the AI Agent provider, or the underlying LLM provider for purposes other than just the AI Agent. The customer could also be subjected to biased profiling.	• Loss of business resulting from incorrect customer assumptions. • Negative reputation as a result of biased profiling. • Your company's profile may be leaked, used, or sold by the AI Agent provider, or the underlying LLM provider.

Understanding the output risks

Risks associated with the content output of the AI Agent.

Risk	Definition	Risk to your customer	Risk to your company
Inaccurate content	The generated output may be factually incorrect, misleading, or even nonsensical.	The customer receives incorrect or unreliable information, resulting in poor decision-making or miscommunication.	Decreased CSAT, and having to deal with the ramifications of the misinformation — this could be as simple as offering a discount to apologize for a misunderstanding about shipping times, or making it right for a customer who lost a sizeable amount of money in an incorrect bank transfer.
Inappropriate or offensive content	The generated output may be inappropriate or offensive, such as perpetuating stereotypes, discrimination, or the spread of harmful ideas.	The customer is on the receiving end of the inappropriate or offensive content, which may cause emotional distress.	• Loss of business and negative reputation implications as a result of the offensive content. • Potential legal and financial implications.
Biased content	The generated output may unintentionally reflect and perpetuate existing biases present in the LLM's training data.	The customer may receive unfair treatment, or be marginalized.	Loss of business and negative reputation implications as a result of biased treatment.
Copyright infringement	The generated output may reproduce or closely resemble copyrighted content.	The customer may be considered to be infringing third-party copyright by using infringing content.	Legal, financial, and reputation implications as a result of violation of copyright law, and/or intellectual property rights.



Legal and security questions to ask while assessing different AI-powered customer service automation vendors

A major factor in minimizing the risks associated with AI Agent is choosing the right vendor. You'll want to partner with a company that has guardrails in place that either omit or drastically reduce the security risks of using generative AI for customer service automation.

Here are some questions to ask that will help you make that decision.

Please keep in mind that this toolkit is meant for information purposes only and does not contain any legal advice. You should seek legal advice on how the use of an AI Agent may impact your business.

Questions about the organization

Is the vendor transparent about their security and privacy practices?

- **Why you should ask this:** This allows you to evaluate their security posture and make a more informed decision.
- **What to look for in an answer:** Yes.

Does the vendor have a strong security posture?

- **Why you should ask this:** Your company's data and your customers' data is your responsibility, and you want to make sure you're partnering with a vendor who takes that responsibility as seriously as you do.
- **What to look for in an answer:** Look for vendors with indicators of strong security posture such as a dedicated security team, compliance with applicable privacy regulations (e.g. GDPR), compliance with security control standards (e.g. SOC 2 Type 2), compliance with applicable industry specific laws (e.g. HIPAA), annual penetration testing, and a bug bounty program.

Questions about personal/sensitive information

Does the AI Agent provide functionality that can automatically redact unnecessary sensitive information in customer inputs?

- **Why you should ask this:** You don't want the AI Agent provider, or any integrated systems, to ingest, store, and/or process any sensitive information that may not be necessary to resolving the inquiry.
- **What to look for in an answer:** Yes.

Does the AI Agent provide a secure way to manage sensitive customer information?

- **Why you should ask this:** In cases where sensitive information is necessary for resolution, you might need the AI Agent to manage the information instead of flat out redacting it.
- **What to look for in an answer:** The vendor should be able to provide a mechanism to securely manage sensitive customer information.

Will my customers' data be used to train the underlying LLM model?

- **Why you should ask this:** You don't want your customers' data to be used to train a third party LLM model as it might result in sensitive data exposure (e.g. LLM provider suffers a data leak).
- **What to look for in an answer:** No, any identifiable customer data is redacted prior to being used to train the model.

Questions about the product

Does the AI Agent incorporate mechanisms that ensure generated content is safe, relevant, and accurate?

- **Why you should ask this:** You want replies to be safe, accurate, and relevant to provide a better service for customers and mitigate output risks.
- **What to look for in an answer:** You'll want to make sure that they specifically have built-in functionality at the product level that addresses each requirement.

Where will the AI Agent source information to generate content?

- **Why you should ask this:** Unknown or public sources (e.g. internet) may lead to unsafe, inaccurate, or irrelevant outputs.
- **What to look for in an answer:** You should be able to connect your knowledge base to the AI Agent so it can source information from and asses for output accuracy.

Is there a human-in-the-loop functionality that requires manual verification of generated content?

- **Why you should ask this:** This additional step further ensures that the generated content is safe, accurate, and relevant.
- **What to look for in an answer:** The functionality should be available but not required, you should be able to trust the automatic filters.

How much access to data will the AI Agent need from other tools in my tech stack to work?

- **Why you should ask this:** It's normal for an AI Agent to need access to backend systems, especially as you start automating more complex resolutions, but you should be able to minimize access to necessary systems only or manage the level of access (e.g. read-only instead of full admin)
- **What to look for in an answer:** You should be able to control which software the tool can access and as well as the level of access it has

Is the vendor able to accommodate regional requirements for data hosting?

- **Why you should ask this:** You may have obligations to your customers to ensure data remains within a specific region and adheres to data sovereignty requirements
- **What to look for in an answer:** Yes

Does the product team adhere to the principles of secure by design, secure by default, privacy by design, and privacy by default?

- **Why you should ask this:** This demonstrates a commitment to incorporating security and privacy into the design of the product — as opposed to an afterthought, or not at all
- **What to look for in an answer:** Yes to all

Questions about the underlying or foundation LLMs

What LLMs does the tool use in providing generative AI features?

- **Why you should ask this:** Knowing what LLMs the tool leverages can help you form a better understanding of potential security risks that may be associated with any specific one
- **What to look for in an answer:** The vendor should be able to provide a satisfactory answer

How does the vendor select what LLMs to incorporate into its services?

- **Why you should ask this:** The vendor should have a list of criteria they use to assess the compatibility of LLMs with the desired use case
- **What to look for in an answer:** The criteria should include reliability, security and data protections, and performance on delivering automated resolutions

Has the underlying LLM provider experienced any notable data leaks or breaches?

Why you should ask this: This is an indication of the maturity of the vendor's security posture

What to look for in an answer: Ideally no, but if the answer is yes, then see how they handled the leak or breach — anyone is at risk of getting hacked, but how an organization responds to and learns from a breach can help you assess their posture

Does the vendor notify customers if new LLMs are incorporated into the AI Agent?

- **Why you should ask this:** LLM providers who process personal data on behalf of the vendor are considered third-party subprocessors under the GDPR and vendors are required to notify customers in advance of using any new subprocessor
- **What to look for in an answer:** Yes

Partner with the pros

We've helped over 350 innovative customer service leaders deliver transformative results using AI. If you're looking for guidance, our consultants are ready to chat about where this technology can fit into your company and how it can help you transform your customer service organization.

Schedule a call

Ada is an AI-powered customer service automation platform on a mission to make customer service extraordinary for everyone. Ada envisions a world where every customer interaction is resolved by AI, and makes it easy for businesses to automatically resolve the greatest number of customer service conversations — across channels and languages — with the least amount of effort.

Since 2016, Ada has powered more than 4 billion automated customer interactions for brands like Meta, Verizon, AirAsia, Yeti, and Square. Born in Toronto, Ada serves companies and their customers worldwide.

T +1 (855) 232-7593

E hello@ada.support

@Ada_CX

facebook.com/getadasupport

linkedin.com/company/ada-cx

youtube.com/@ada_cx

ada.cx