PixieBrix

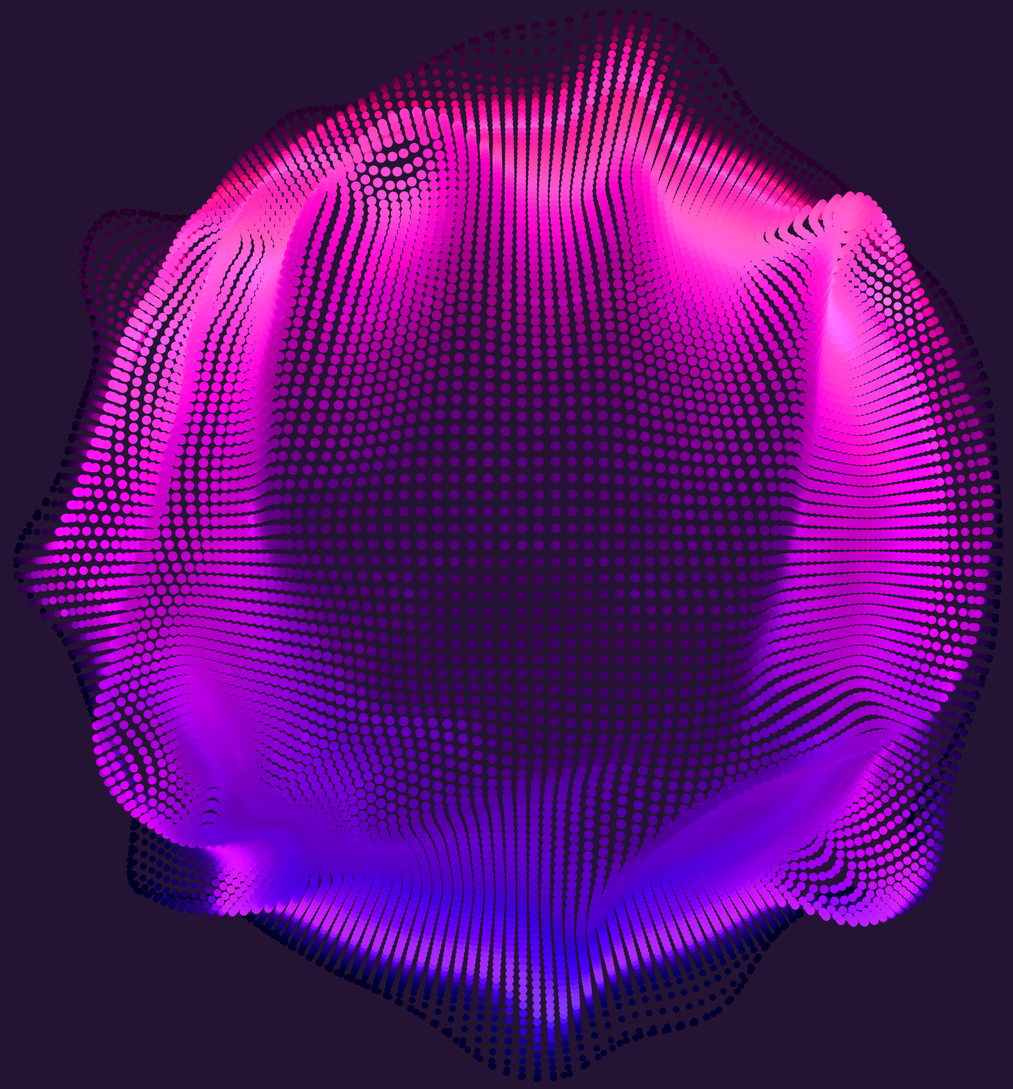

Adopting AI in Customer Support

PixieBrix's State of AI in Customer Support

2025 Report

pixiebrix.com/reports/ai-customer-support-report

Table of Contents

If you're a CS leader looking to deepen your understanding of the AI landscape, this guide is for you. It's not a list of "100 GenAI use cases," but instead a set of frameworks and questions to help you understand, navigate, and successfully adopt today's popular AI tools.

Key Insights

KPI Focus

It's tempting to look at a new technology and see how it could fit. However, you should first pinpoint an internal KPI, and assess technology solutions based on impact. A more prescriptive approach leads to better business results, without getting swept away in possibilities.

Solution Sweet Spots

Given high feature overlap among software providers, it's important to consider solution sweet spots. For example, ticketing systems, chatbots, and copilots are interconnected but require different capabilities and technology. Your unique support needs, KPIs, and resources should guide whether an add-on solution or best of breed makes sense.

Managing Risk

AI pilots fail for many reasons; we often cite context engineering, user experience, technology selection, and business alignment. To save yourself a headache in the long run, invest time and money in a technical POC or sponsored pilot to de-risk any solution.

IT Buy-In

IT can be a powerful ally in your AI journey. IT executives have the purview, resources, and mandate to assess AI solutions from a company-wide perspective. When CS becomes a partner or launchpad for broader AI platform investments, it opens the door to new options, more control, and better economics.

Hype vs. Reality

Gen AI has the potential to shift business models and reshape industries. How soon will AI agents take over? That's the wrong question. If you're a CS leader, you know AI is ubiquitous, changing what skills your team needs, how customers interact with your brand, and more every day. The prudent strategy is early adoption vs. wait and see.

What You Need to Know About Gen AI (in 200 Words)

By now you've used generative AI popularized by OpenAI's ChatGPT. You've asked it a question or given it instructions ("prompted it"), and seen it write poetry, code, emails, and more.

In Customer Support, Gen AI offers three broad capabilities:

1. Content generation (e.g., responding to emails)
2. Search/retrieval (e.g., suggesting the right KB article)
3. Task automation (e.g., creating a bug ticket)

The Rise of Agentic AI

You can now give gen AI access to software tools and company data, and with the right orchestration, the AI system can succeed at ambitious tasks that would normally require humans.

This breakthrough is called "Agentic AI" and these are the "AI Agents" you read about. To clarify, this is task automation within a controlled system, not artificial general intelligence (AGI).

That said, there's little doubt that agentic AI will eat into call center, accounting, software development, legal services, and beyond.

30,000 Foot View: Where AI is Shaping CS

This report focuses on four categories where we see the most AI disruption: Customer Service Platforms, Chatbots, Copilots, and Enterprise Knowledge.

The category segmentation emphasizes tradeoffs around core focus. We list 1-3 popular solutions in each segment to illustrate the point for market context, not as a vendor ranking or comprehensive market map.

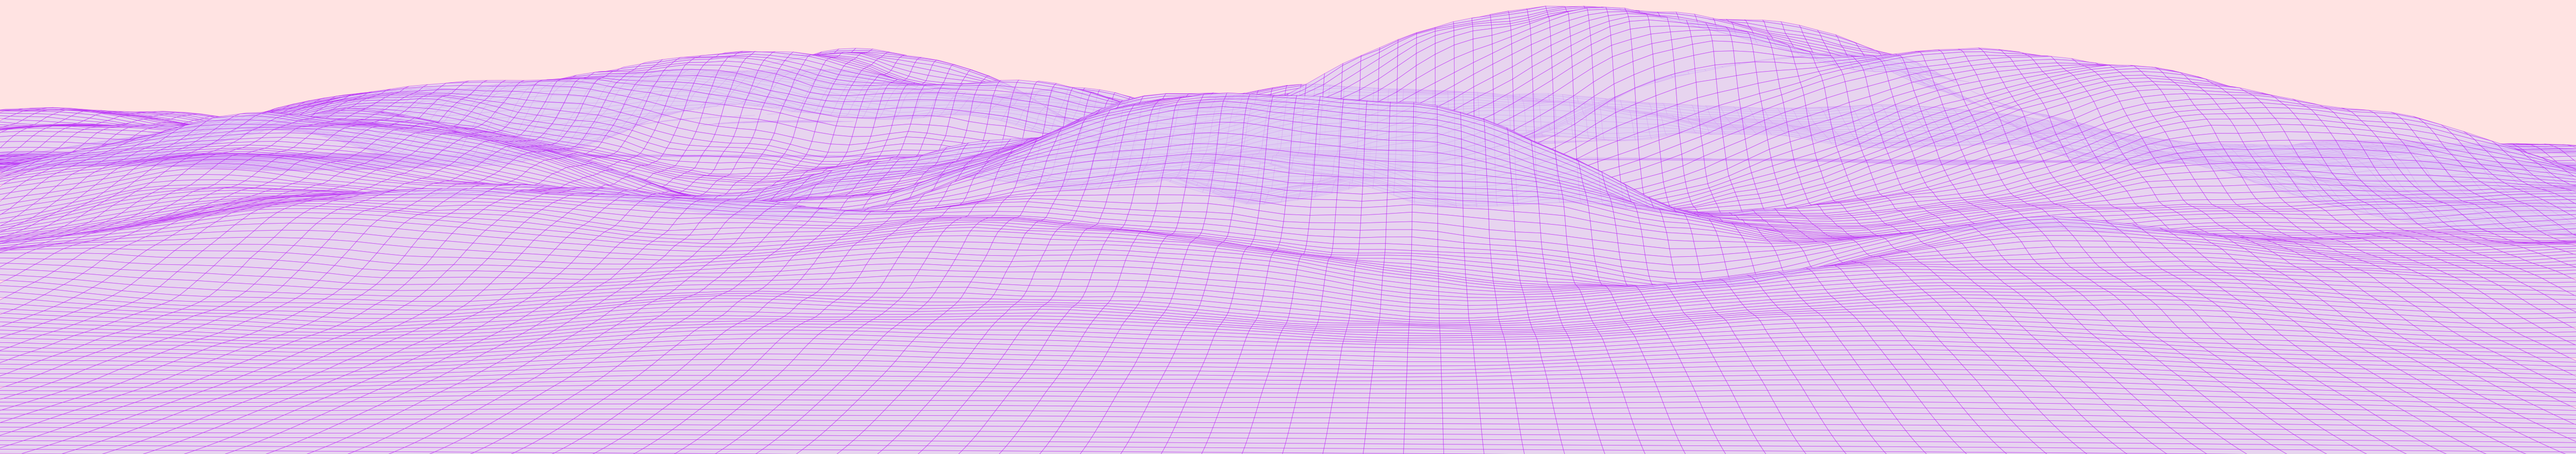

Category	Segment	Scale	Summary
Customer Service Platforms	**Ticketing Systems** Intercom Zendesk	Intercom = VC-backed Zendesk = $10B (PE)	Support-focused To resolve customer issues
	CRMs HubSpot (Service Hub) Salesforce (Service Cloud)	HUBS = $30B CRM = $250B	Customer relationship-focused To manage the customer lifecycle from marketing, sales, support, and beyond
	CCaaS Platforms Verint Five9 Genesys	VRNT = $1B FIVN = $2B Genesys = $20B (PE)	Contact center-focused • Intelligent routing • Workforce optimization • Conversational AI Omnichannel: voice, email, chat, social
Chatbots *Customer facing AI bot*	**Answer Bots** kore.ai Yellow.ai	VC-backed	First generation rules-based chatbots Have evolved; marketed as "agentic platforms"
	DIY Azure AI Amazon Q	MSFT = $3.5T AMZN = $2.5T	IDE for building agentic chatbots ("agents") Designed to work w/ complimentary services (MSFT Copilot Studio) Geared toward technical users
	Add-ON to CSPs Intercom (Fin) Genesys (Virtual Agent) Salesforce (Agentforce)	Intercom = VC-backed Genesys = $20B (PE) CRM = $250B	Easily integrated with data in the native platform
	White Glove - Agentic Chatbot Decagon Forethought Sierra	VC-backed	Full service platform for rolling out chatbots powered by agentic AI ("agents") Geared toward self-serve question answering i.e. deflection

Figure X: Categories and Segments

Category	Segment	Scale	Summary
Copilots *Employee facing AI assistant*	**Native Copilots** Microsoft Copilot Google Gemini	MSFT = $3.5T GOOG = $2T	Great for personal productivity i.e. writing, brainstorming, coding, & potentially search Can be customized for CS use cases Designed to integrate within the native ecosystem
	Copilot Platforms PixieBrix	VC-backed	Customizable AI assistant geared toward high-touch customer support Integrates and embeds inside diverse apps (not tied to one ecosystem) Spikes on extensibility, fine-grained permissions, and UX
	Add-On to Chatbots Forethought (Assist) Decagon (Agent Assist)	VC-backed	Like customer-facing chatbots, but with access to internal company data
	Agent Assist Uniphore (Real-time Guidance Agent)	VC-backed	Real-time, in-call agent assistance and guidance using conversational AI Omnichannel, 360 customer view
Enterprise Knowledge	**White Glove - Enterprise Search** Glean Algolia	VC-backed	First generation rules-based chatbots Have evolved; marketed as "agentic platforms"
	Knowledge Management Confluence Guru ClickUp	TEAM = $40B Guru/ClickUp = VC-backed	IDE for building agentic chatbots ("agents") Designed to work with complimentary services (e.g. MSFT Copilot Studio) Geared toward technical users
	Add-ON to CSPs Salesforce Knowledge Base Zendesk Knowledge Base Intercom Knowledge Base	CRM = $250B Zendesk = $10B (PE) Intercom = VC-backed	Easily integrated with data in the native platform
	Learning Management Sana Labs	VC-backed	Create and organize training materials "Digital classroom"

Figure X: Categories and Segments

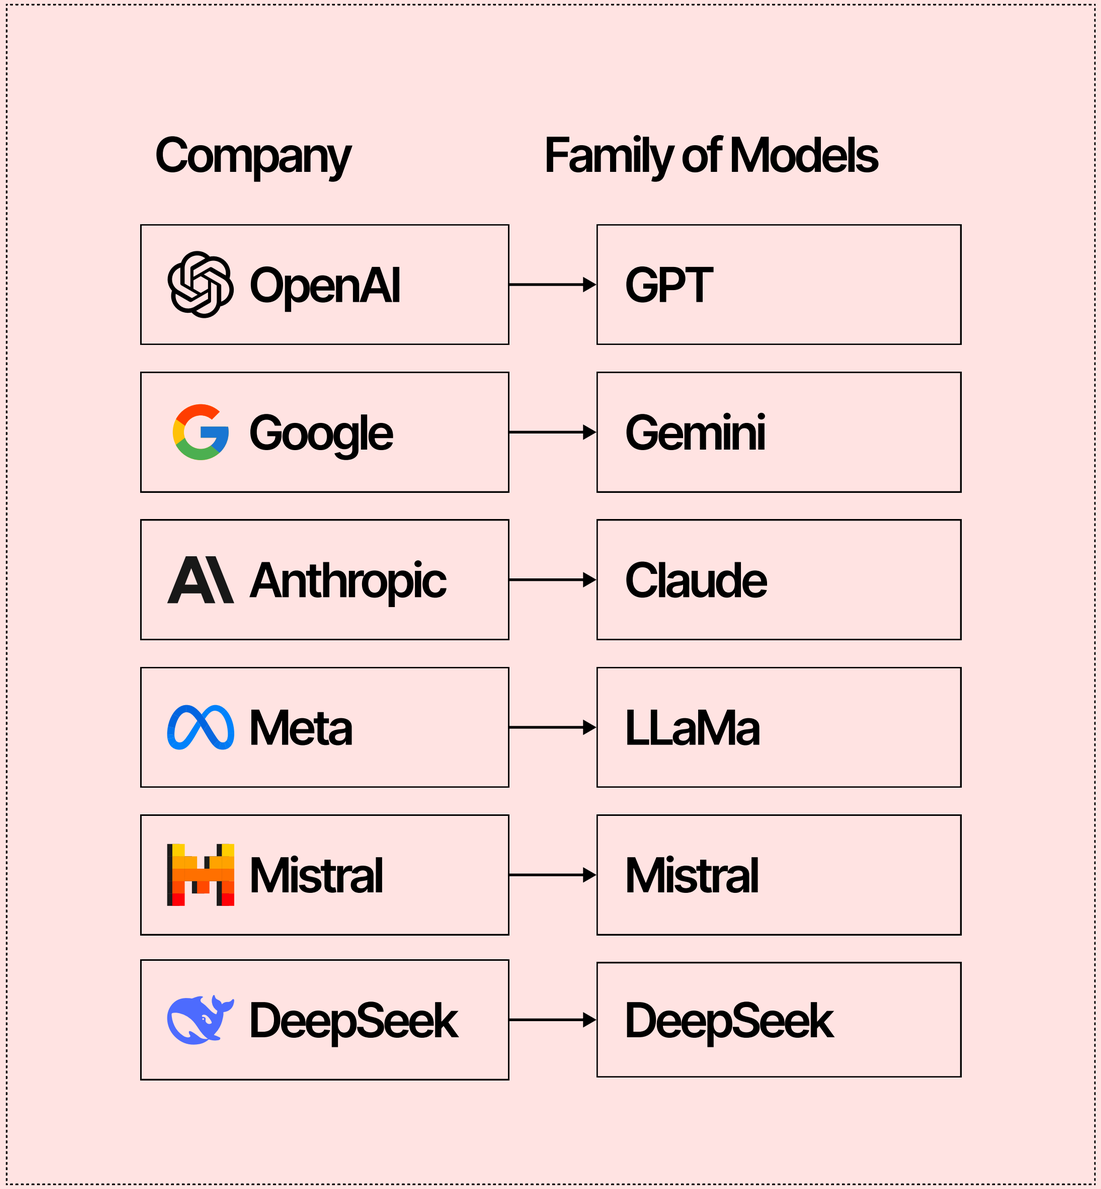

Foundation Models

By the way, most software applications use the following models in some way "under the hood."

Note: specific models have defined specs with names like "GPT-4o mini."

As you can see, the landscape is a mix of established software providers and disruptors, and solutions differentiate on "sweet spots" and can cash in with "add-on" capabilities. A feature of one product can be an entire separate company altogether.

So how do you navigate it?

Five Steps for Adopting AI

1. **Clearly define your problem and KPIs**
   "What am I trying to solve, and how would I measure a solution?"
2. **List key criteria and challenges that solution needs to address**
   "What am I trying to solve, and how would I measure a solution?"
3. **Consider tradeoffs based on solution archetypes**
   "Does this solution make sense for my business?"
4. **Gauge the cost of moving forward**
   "What's my full investment?"
5. **De-risk your decision**
   "How do I get proof before I go all in?"

Step 1: Clearly define your problem and KPIs

“What am I trying to solve, and how would I measure a solution?”

Start with the Pain

Before evaluating solutions, start with the real operational challenges support leaders face daily:

- Training is long and ineffective, hurting quality
- Tickets keep getting escalated to L2/L3, driving up costs
- Agent responses lack consistency, impacting CSAT
- Agents get lost navigating tools, creating frustration and low morale
- My knowledge base has gaps and we can’t keep up, driving up AHT

Pinpoint the KPI

Across companies, we consistently see three primary goals emerge when adopting AI in support:

Goal	Tactic	KPI
Reduce Ticket Volume	Enable customers to directly self-serve without talking to a human support agent	Deflection Rate
Reduce Escalation	Empower frontline or L1 human support agents to be more self-sufficient	Escalation Rate
Reduce Handle Time	Empower CS teams to solve inquiries more efficiently	AHT or MTTR

The three KPIs mentioned above all tie into cost of service, which makes ROI clear. Cost of service is a key input to the AI business case. That's not to say CS leaders aren't intensely focused on quality and customer satisfaction, and any solution should drive these metrics, too. We know by now that CS shouldn't be viewed purely as a cost center.

We've included some KPI benchmarks in the appendix for your reference.

Narrow Your Research

Different solutions are geared toward different primary KPIs. For the example:

KPI	Category
Deflection Rate	Chatbot
Escalation Rate	Copilot
AHT or MTTR	Copilot

Keep in mind, KPIs don't tell the whole story. Just ask your colleagues about deflection:

Do all chatbot interactions represent true deflection?

Does 40% deflection mean ticket volumes went down by 40%?

Is your team still underwater?

Did you avoid customer blowback like Klarna?

Step 2: List key criteria & challenges that solution needs to address

“What makes my situation unique? What am I looking for?”

We chatted with Jen, a CS leader who oversees 200+ support engineers at a global SaaS company. Jen's AI journey was grounded in a clear KPI: mean time to resolution (MTTR). Jen summarized her unique challenges and solution criteria:

Bio

"We're not trying to replace people - we're trying to scale the brilliance of our top agents and reduce noise for the rest."

Jen oversees a globally distributed team of 200+ support engineers across multiple time zones. Her day-to-day centers around maintaining high-quality support experiences while reducing escalations and resolution times. Jen comes from a technical background but leads with empathy - she's passionate about enabling agents with the tools they need to handle complex, high-urgency cases efficiently.

She's not new to AI, but she's pragmatic. Her interest lies in tools that help reduce context switching, auto-summarize cases, and streamline L1–L2 handoffs - without requiring months of onboarding or expensive change management.

Quick Facts

Role	VP, Customer Support
Company	Global B2B SaaS Company
Team Size	200+ support engineers
Region	US, with global support coverage
Primary Channel	Email and chat (via Zendesk)
Tech Stack	Zendesk, Salesforce, Confluence, Slack
KPI	MTTR (Mean Time to Resolution)

Goals

- Shorten MTTR across support tiers
- Reduce reliance on expensive L2/L3 escalation
- Improve case triage and auto-summarization
- Equip agents with AI-powered suggestions that feel reliable

Key Challenges

- Cases often involve 3rd-party dependencies and bugs
- Agents work with fragmented, incomplete sources of truth
- Access levels and permissions vary across systems
- Escalations are expensive and time-consuming
- Agent skill levels vary widely - making quality inconsistent

Solution Criteria

Category	What They're Looking For
Tool	Copilot
Channels	Must support Zendesk (email/chat)
Integrations	Native or light-plug Zendesk integration
Priority Features	Summarization, suggestions, next-step recs, draft replies
Differentiators	Zendesk, Salesforce, Confluence, Slack

Consider the dimensions below when summarizing your company's unique challenges.

About You: Key Challenges

Dimension	Considerations	Implications
The Nature of Your Product or Service	B2C vs. B2B How technical or complex? How many different offerings? What price points?	More complexity = more human touch; emphasis on collaboration, integrations, customization, UX
Your Support	**Modes of communication** • Channel mix (email, chat, voice) • Language requirements **Scale Complexity** • Hiring / onboarding frequency • Team organization; in-house vs. BPO **Human Capital** • Number of human agents • Avg. compensation	Binary requirements to be addressed by scope of solution Scale complexity necessitates end user simplicity and strong control around end user permissions Puts cost base and ROI into perspective
Existing Investments	**Technology** • What core apps? (ticketing, KB, WFM, chat, portals) • Vendor concentration? (Google, MSFT, Zendesk) • Web based and desktop? **Content** • Is your KB organized? How big are the gaps? • Other critical info sources? (docs, tickets) • Content format? (text, processes, images)	Key to questions around integrations and compatibility Potential switching costs Hidden cost to making solutions Potential root cause of failure
IT Alignment	**IT Buy-In** • Are you partnering with an IT executive? • Is there a broader AI initiative? • Do you have access to technical resources? **Security** • What's required (GDPR, SOC 2, PCI, etc.) • How varied is data access? • Data sync vs. zero copy?	Key to questions around integrations and compatibility Potential switching costs Hidden cost to making solutions Potential root cause of failure

Now let's look at the key dimensions to evaluate AI solutions:

Dimension	Considerations
AI	LLM support: open vs. closed source Bring your own model? Response accuracy (fine tuning, context engineering) Agentic (MCP, orchestration, guardrails) Observable, auditable
Integrations	Out-of-the-box integrations with CS tools and data sources Low-code integration builder Snapshot vs. live data connection? Knowledge graph vs. zero copy?
Compatibility	Is data migration required What channels are supported Web vs. desktop
User Experience	Customizable, Embedded, Simple
Setup & Maintenace	Are technical resources needed Time to value
Security	Authentication: Single Sign-On (SSO) Permissions: fine grained by user/group Data flows (PII, etc.) Security standards (PCI, etc.)

Step 3: Consider tradeoffs based on solution archetypes

“Does this solution make sense for my business?”

Each solution comes with tradeoffs around time, budget, and control. Consider the following:

1. Implementation complexity
2. Operational maintenance
3. Vendor lock-In

Implementation Complexity

To what extent does the solution require technical integration, resourcing, or bespoke configuration?

Consideration	Strategic Implementation
Implementation Timeframe	Lengthy setup timelines can delay time-to-value and strain internal technical resources.
Technical Resourcing	Solutions requiring engineering input may hinder agility, particularly in lean support teams.

Operational Maintenance

What internal resources are required to maintain and optimize the AI assistant over time?

Consideration	Strategic Implementation
Workflow Updates	High-maintenance tools may require continual training or logic updates (e.g. intent management).
Content Drift	As support documentation and business logic evolve, systems must adapt seamlessly to remain effective.

Vendor Lock-In

What are the switching costs if the solution no longer meets business needs?

Consideration	Strategic Implementation
Ecosystem Dependency	Tools embedded in proprietary stacks (e.g., Microsoft, Salesforce) can be difficult to migrate from.
Workflow Portability	Platforms that don't support export or modular workflows increase long-term switching costs.
Data Ownership	Ensure data can be exported in usable formats (e.g., ticket history, AI logs, prompt configurations).

Solution Archetypes

Based on several factors including time, budget, and control, we've created a few solution archetypes. This is not a comprehensive list but rather a sample to illustrate tradeoffs.

Add-on Approach

CS Platform

E.g. Salesforce Agentforce

Considerations

- ✓ Does the solution meet my performance criteria?
- ✓ Does it integrate outside the platform?
- ✓ Will it address my needs as they change over time?
- ✓ Am I ok with how they do permissions?
- ✓ Am I getting too "locked in?"

Budget

Your Time

DIY Approach

One Ecosystem

E.g. Azure AI Services

Considerations

- ✓ Do the "build vs. buy" tradeoffs make sense?
- ✓ Is this where I should focus technical resources?
- ✓ Am I okay with slower time to value / execution risk?
- ✓ Am I getting a great deal?

Budget

Your Time

White Glove Approach

Chatbot or Search

E.g. Glean

Considerations

- ✓ Does the performance justify the price?
- ✓ How hard is it to reverse this decision?
- ✓ Am I okay with permissions and data flow?
- ✓ Can I become self-sufficient over time?

Budget

Your Time

Integrator Approach

Copilot Platform

E.g. PixieBrix

Considerations

- ✓ Does an "integrator" approach make sense?
- ✓ How much do I value flexibility / optionality?
- ✓ Do I have technical resources or can I partner with IT?

Budget

Your Time

Step 4: Gauge the cost of moving forward

“What’s my full investment?”

AI solutions have moved toward usage and outcome based pricing. While this can create alignment, it can also lead to confusion. This section can help you dissect the pricing components:

Pricing Models

How does pricing model align with the organization's usage patterns, scale, and definition of value?

Model	Description	Strategic Considerations
Per Seat	Fixed monthly fee per user	Predictable budgeting May not reflect actual tool utilization
Per Ticket Volume	Fees based on number of support tickets processed or touched	Do all interactions qualify as tickets? What about volatility / spikes?
Per Resolution	Pricing based on deflected or resolved tickets	Requires a clearly defined and mutually agreed-upon success metric
Hybrid	A blend of the above models	Combination of questions above

Total Cost of Ownership (TCO)

As discussed, you pay in more ways than one. Consider the following categories when evaluating total cost of ownership.

Vendor Flexibility

Ease of integration, switching costs, interoperability

Scalability

Predictability of cost with growth (agents, tickets, geographies)

Setup

Time to deploy, onboarding effort, implementation support

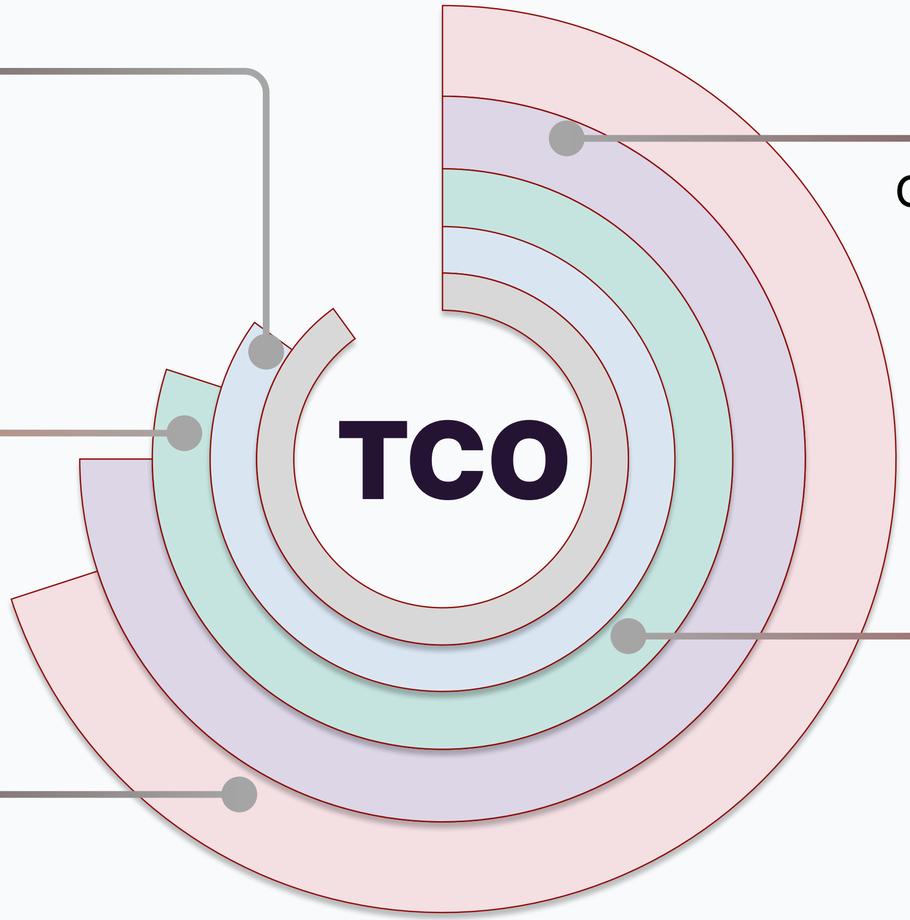

Licensing

Core pricing model, flexibility of contract terms, usage tiers

Ongoing Operations

Maintenance burden, training costs, admin ownership

Pricing Model Scenario - Intercom Fin (Real Data)

Context:

- You're evaluating Intercom's Fin AI Agent for deflecting Tier-1 support tickets.
- Your team runs Zendesk (or another helpdesk) - not Intercom's platform.
- Fin charges $0.99/resolution, with a minimum of 50 resolutions/month, and no per-seat fees.
- Fin requires underlying helpdesk seats if using Intercom's native plan: $29–$132 per seat/mo (billed annually).

Category	Details	Fin (Intercom) Scenario	Strategic Implications
Pricing Model	Pay per resolved conversation (not seats)	$0.99 per resolution, minimum 50/month	Outcome-aligned pricing. Great for scaling, but subject to monthly volume volatility.
Setup	Time to deploy, onboarding effort	Requires integration with Intercom's stack. Migration from Zendesk/other platforms may take 4–8 weeks.	Slower time to value if not already in Intercom ecosystem. May increase IT lift.
Licensing	Core pricing model, contract structure	No per-seat charge; resolution-only billing. Annual commitment typically required.	Cost scales with usage, not team size. Budgeting becomes variable, not fixed.

Category	Details	Fin (Intercom) Scenario	Strategic Implications
Ongoing Operations	Maintenance, content tuning, internal ownership	Requires frequent tuning of KB, resolution mapping, and escalations. Intercom controls many aspects.	Internal CS ops or vendor management overhead needed. Agent-AI collaboration requires clear routing.
Scalability	Growth across tickets, agents, geographies	Easily expands across agents with no extra cost. Cost tied to deflection/resolution volume.	Predictable scale if performance is stable, but price can spike during volume surges.
Vendor Flexibility	Switching cost, ecosystem lock-in, interoperability	Only runs within Intercom stack. Difficult to decouple from broader Intercom platform.	High switching cost. Limited flexibility if you use Zendesk, Salesforce, or other support platforms.

Assumptions	Estimate
Ticket volume	10K/mo
Estimated resolution rate (AI handles)	60% → 6,000 Fin resolutions
Monthly Fin Cost	6,000 x $0.99 = $5,940
Baseline spend (minimum)	50 resolutions → $49 + admin cost = negligible

Let's Explore Different Scenarios

Scenario 1

Low Volume / Pilot Phase

Est. Monthly Cost: $248

Assumptions

500 tickets/month

50% AI resolution rate (250 resolutions)

$0.99 per resolution

Budget

Your Time

Best for testing AI in production without high upfront costs.

Spend remains predictable but may not fully showcase ROI potential.

Easy to pause or scale up.

Scenario 2

Mid Volume / Moderate Deflection

Est. Monthly Cost: $2,722

Assumptions

5,000 tickets/month

55% AI resolution rate (2,750 resolutions)

Budget

Your Time

Delivers measurable agent capacity relief.

Budget variability starts to increase with ticket spikes.

Still cheaper than hiring additional Tier-1 agents.

Scenario 3

High Volume / Enterprise Deployment

Est. Monthly Cost: $5,940

Assumptions

10,000 tickets/month

60% AI resolution rate (6,000 resolutions)

Budget

Your Time

Maximum operational leverage - significant deflection from Tier-1 queue.

Monthly cost fluctuates heavily with ticket surges.

Requires strong KB governance and AI performance monitoring.

Step 5: De-risk your decision

“How do I get proof before I go all in?”

These iceberg charts are great because they really illustrate the true cost of making AI solutions work. The “tip of the iceberg” is the AI demo, but then there’s everything below the surface.

Therefore, we recommend a robust evaluation before going “all in” on any solution.

The Depths of AI

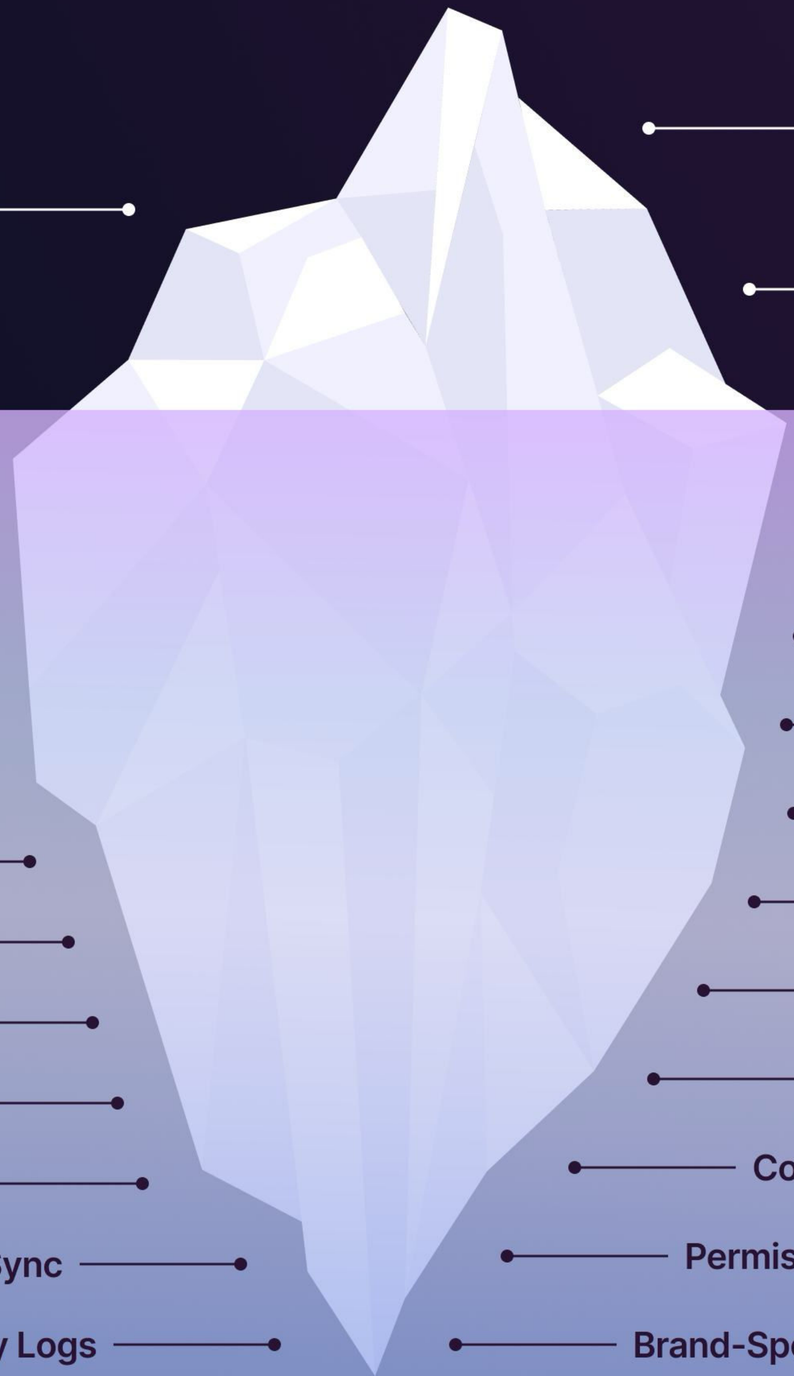

Generative Language Model

Retrieval-Augmented Generation

App or Workflow Interaction

Multi-Step Workflow Logic

Contextual Guardrails

Sensitive Data Recognition & Masking

Custom Model Training

Stability & Behavior Regression Checks

Team-Based Quality Assurance

Content-Driven Branching

Deployment Workflow Controls

Cross-Language Communication Capability

Policy Enforcement

Structured Data Sync

Automation Activity Logs

AI Decision Trails

Input Validation

Performance & Behavior Tracking

Omnichannel Communication Support

Automation Change History

Live Agent Escalation

Version Control for AI Systems

Agent Behavior Test Environments

Concurrent Processing for Faster Response

Permissioned Access Layers

Brand-Specific Terminology Control

Three Ways to Evaluate Solutions

Use below framework to gradually increase confidence and de-risk decision. Note that these are not necessarily mutually exclusive.

Evaluation Method		What It Tells You
Free Trial	→	Can I get it to run on my data?
Technical Proof of Concept (POC)	→	Does the core tech work with my stack, data, and workflows?
Pilot / Proof of Value (POV)	→	Does it drive results at scale + get buy-in from frontline users?

Pros & Cons

Evaluation Stage	Pros	Cons
Free Trial	Fast to start No commitment Good for small teams	Limited vendor support Easy to misconfigure May not reflect real usage
Technical POC	Validates stack compatibility Verifies core features	Doesn’t show agent UX May miss operational edge cases Not a true “battle test”
Pilot / POV	Closest to real deployment Lets teams trial actual workflows	Highest effort / takes time + alignment Must track KPIs and agent feedback rigorously

Conclusion

The rapid pace of AI innovation can be daunting. However, you don't have to be an AI expert to bring on AI solutions - you just have to know what questions to ask. We hope this research helps.

To that end, let's reinforce three key takeaways:

1. Your KPI should dictate which category of software solutions to research.
2. Within a given category, there are various considerations around not just the AI, but also the integrations, compatibility, user experience, setup & maintenance, and security.
3. The more you commit and invest up front, the better results you'll see in the long run.

Appendix

Goal	KPI	Typical Baseline (Pre-AI)	Benchmarks with Assist AI
Reduce Ticket Volume	Deflection Rate	10–25%	25–45% (Retail) 20–35% (SaaS)
Reduce Escalation	Escalation Rate	30–40%	15-25%
Improve Handle Time	Avg. Handle Time (AHT) / MTTR	AHT: 5–15 mins	15–30% faster
Improve Support Quality	CSAT / NPS	CSAT: 75–85%	CSAT: + 5–15 pts
Improve Agent Satisfaction	Agent NPS	n/a or < 30	+ 20–40 post-pilot