Achieving a 64% Self-Service Resolution Rate Through a Proactive AI Agent.

Problem Statement
Support teams using our live chat platform are overwhelmed by high volumes of repetitive inquiries. This inflates operational costs and prevents skilled agents from focusing on complex issues, ultimately leading to slow response times and a frustrating experience for their end-users.

Role
Lead Product Designer

Key Achievement: Improved chatbot resolution rates to 64% from 44%

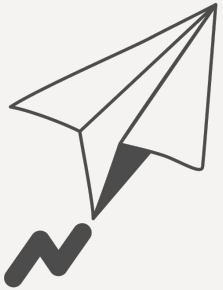

OVERVIEW

HubSpot (Conversational Assistants)

Challenge

Navigating Technical Uncertainty: A primary challenge was designing the agent's conversational flows and management UI in parallel with engineering's exploration of a new AI framework. This required a deeply **collaborative and iterative** process to ensure the user experience could adapt to emerging technical capabilities and constraints.

Team

Sr. Product Designer (Myself), Sr. Product Manager, BE engineers, FE engineers, Content Designer

Scope/Constraints

To launch the MVP on a tight timeline, the V1 AI agent was **intentionally trained on a focused set** of our most critical support documents. The primary **design consideration was to create a valuable and trustworthy user experience**, even when the agent's capabilities were narrow, by managing user expectations and providing seamless escalation paths to a human agent.

Discovery

QUANTITATIVE ANALYSIS

How are the current automated flows performing?

Collaboration: PM

To establish a **data-driven foundation** for the project, I conducted a quantitative analysis using Amplitude. My goal was to understand the **performance and limitations of our existing automation features**.

I specifically investigated which customer accounts were using rules-based chatbots, and **how often these flows were handed off to human agents**. This would be a key indicator to how effective the bots were in resolving support tickets.

The insights gathered were that **rules-based bots were resolving tickets at a 44% rate**. Far below customers expectations. This increased the load on their front office support staff and **led to our aligned success metric**.

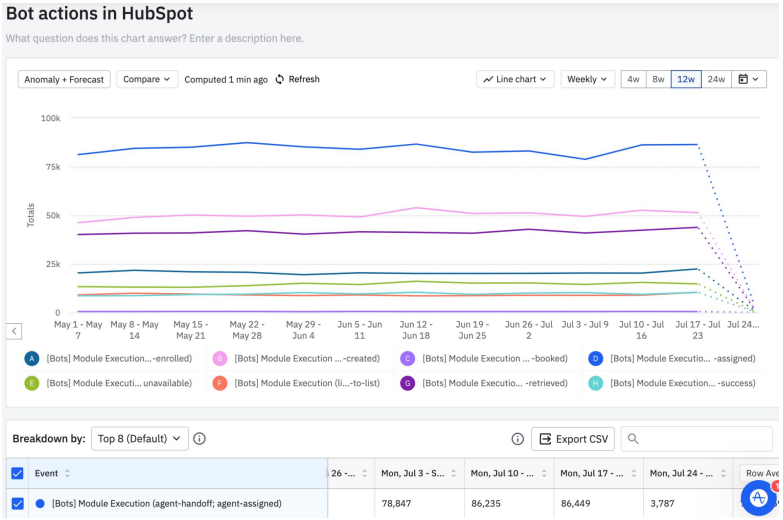

56%

Rules-based bot flows are handed off to human agents.

QUALITATIVE ANALYSIS

How can HubSpot improve customer's workflows?

Collaboration: PM

While the Amplitude data revealed *what* was happening—that our current rules-based bots had **high adoption but low effectiveness**—we needed to understand the *why*. To uncover the story behind the numbers, I partnered with my Product Manager to launch a qualitative research study.

Using our data analysis, we recruited 10 customers who fit our target profile: **high chat volumes and an urgent need for better automation**. Through these in-depth interviews, our goal was to map their current support workflows, pinpoint specific frustrations, and **validate our core hypotheses for an AI agent**.

Goals

1. How can AI bring value to our customer's business?
2. What are the key use cases?
3. What are customer's key pain points with rule-based bots?
4. What are customers' knowledge of AI and willingness to use it?
5. How can success be measured with an AI integration?

RESEARCH SYNTHESIS

What are customers saying?

Collaboration: PM

To transform our raw data into a cohesive strategy, I worked to synthesis our quantitative, qualitative, and competitive research streams. To accelerate the analysis of our qualitative data, I leveraged **Google's LLM Notebook** to process the PRD, research plan, and interview transcripts to rapidly identify recurring patterns, pain points, and user quotes.

In close collaboration with my Product Manager, we then **combined these AI-driven insights with our quantitative metrics and competitive findings** to build a comprehensive, evidence-based understanding of the core problem to be solved.

Personalization

“My chatbot is good at handling basic questions we have accounted for. However, it would be great if the conversation could be more fluid”

Trust

“I’m not sure about the AI hype. How can we ensure it is answering correctly?”

Seamless Handoff

“An AI solution could help our support reps, however we would need a trigger that allows our users to reach a live agent at any time.

PERSONAS

Who are our core customers?

Collaboration: PM

Front Office Admins

Mid-Size business specialist that focuses on creating, editing bots and also responding to chats - 25-40/day. Edits chatflows once every few weeks.

Needs

- To address product questions coming from visitors in a timely manner
- To close leads faster by triaging incoming conversations and getting them routed to the right person
- Bots to automate more conversations

End Users

Customer's visitors who are engaging in conversations through HubSpot's live chat widget.

Needs

- Getting answers in a timely manner
- Streamlined conversations that are not stuck in a bot loop
- Getting honest and helpful answers that mirror what a human interactions

COMPETITIVE ANALYSIS

What are best practices, trends, and opportunities in the market?

To carve out a unique position for HubSpot in a crowded market, I conducted a comprehensive competitive analysis. I **audited the AI chat solutions** of key competitors, evaluating their end-to-end user experience—from conversational UI patterns to the seamlessness of their human-agent handoff.

This process allowed us to benchmark **best-in-class features** and, more importantly, identify a critical market opportunity to differentiate by offering a more intuitive and powerful solution.

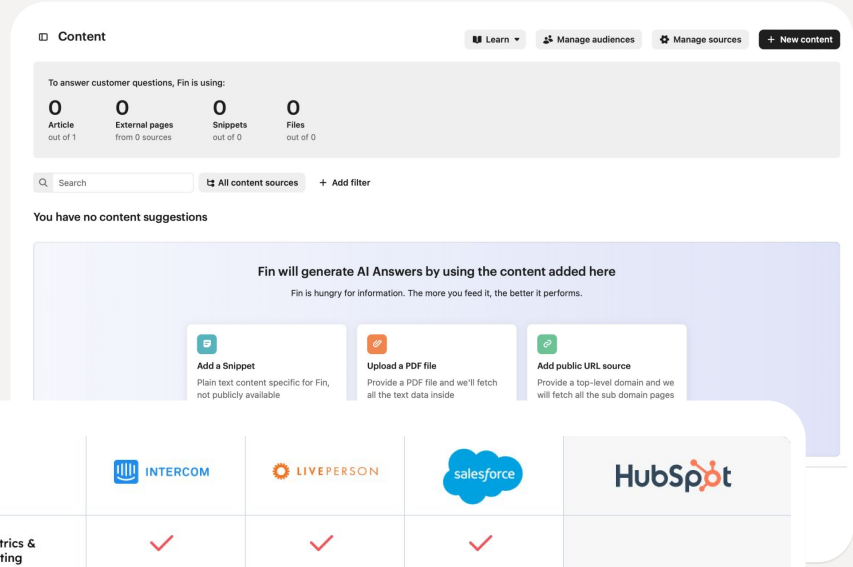

	INTERCOM	LIVEPERSON	salesforce	HubSpot
Chat Metrics & Reporting	✓	✓	✓	
AI Features	✓	✓	✓	MVP
AI Settings Page	✓	✓	✓	MVP
Data Sources (URL)	✓	✓	✓	MVP
Data Sources (Non-URL)	✓	✓	✓	
Bot Builder + AI Integration		✓	✓	
Testing & Preview AI Bots	✓	✓	✓	

DESIGN BRIEF

What are we solving for and how do we measure success?

Collaboration: PM

To translate our research synthesis into a clear, **actionable strategy**, I authored a UX Design Brief.

This document served as our team's North Star, articulating the **core user problem** we were solving, defining our key **design principles**, and outlining a UX vision that directly addressed the needs and pain points uncovered in our research.

This brief also **highlighted the success metric we aligned on from my quantitative analysis**.

Defining Success

Ticket Resolution Rate

Customers will configure an AI Chatbot profile to run across multiple channels to meet their users where they are	**Customers want the ability to run an AI chatbot across multiple channels that feeds into one specific inbox.** (allowing for more targeted experiences for their end users and improved efficiency for their front office conversations)

See more detail on who we're solving for in our CATS Vision Deck.

Goals

Our goal is to create a MVP version of the AI Chatbot to further empower our customers to improve their front office and end user's experiences.

We want to:

1. Scale this product to seamlessly work across all communication channels
2. Simplify the setup and configuration experience
3. Give our users the flexibility, within a controlled manner, to deploy bots that they can trust
4. Educate our customers on how AI Chatbots can work alongside existing products to ultimately help streamline their conversations across the front-office

While keeping these things in mind:

- We want to lean into usage-based pricing and build in-product hooks to drive usage, activation and tier upgrades
- The technology is dynamic and constantly evolving so we will do our best to adapt to evolving customer needs and have a tight feedback loop with users to validate assumptions

PROBLEM STATEMENT

What is our customers core problem?

Collaboration: PM

Support teams using our live chat platform are **overwhelmed by high volumes of repetitive inquiries**. This inflates operational costs and prevents skilled agents from **focusing on complex issues**, ultimately leading to slow response times and a frustrating experience for their end-users.

VISION DECK

How do we scale past V1 and win in the market?

Collaboration: PM

I collaborated with my PM to craft the product vision deck, where I was responsible for **articulating the user-centered aspects** of our strategy. I synthesized our research to define the **core user problems and competitive opportunities**, which laid the groundwork for the vision.

A key contribution of mine was developing the **conceptual UX mockups** that illustrated what the end-state experience would look like, making our **strategic goals tangible and inspiring** for the entire team. This deck was actively updated to reflect new insights as we progressed from private to public beta.

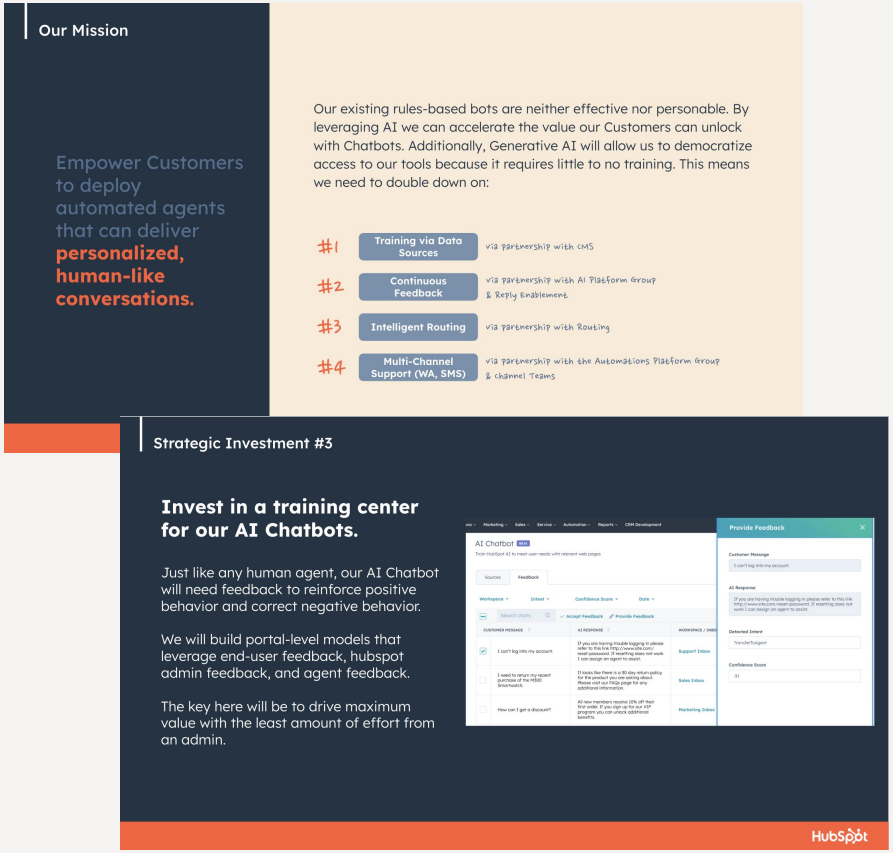

Exploration

WORKSHOP

What elements make up our system today? How can an AI agent fit in?

Collaboration: PM, BE/FE ENG

To create a shared understanding of our ecosystem, I facilitated a **collaborative workshop with key stakeholders** from product and engineering. I developed a detailed systems map which we used as a conversational centerpiece to visually anchor our discussion.

The primary goals were to collectively **identify the relationships of live agents & rules-based bots,** within the current system. Ultimately helping to frame our approach to the AI agent as a personalized experience.

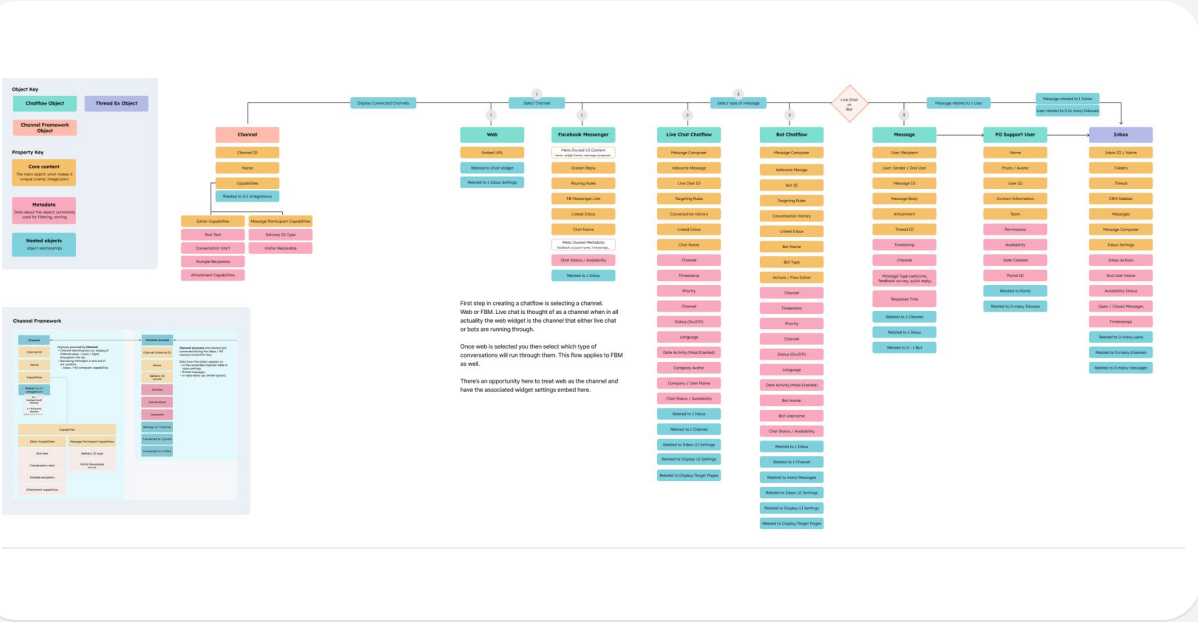

WORKSHOP

What will our customers journey look like?

Collaboration: PM, BE/FE ENG

With a shared understanding of our ecosystem, I facilitated day 2 of the **collaborative workshop centered around user flows.**

The primary goals were to collectively identify **the current system structure** and flows for **handoffs and assigning** conversations.

Identifying opportunities and aligning on how our AI agent could be assigned within our current ecosystem.

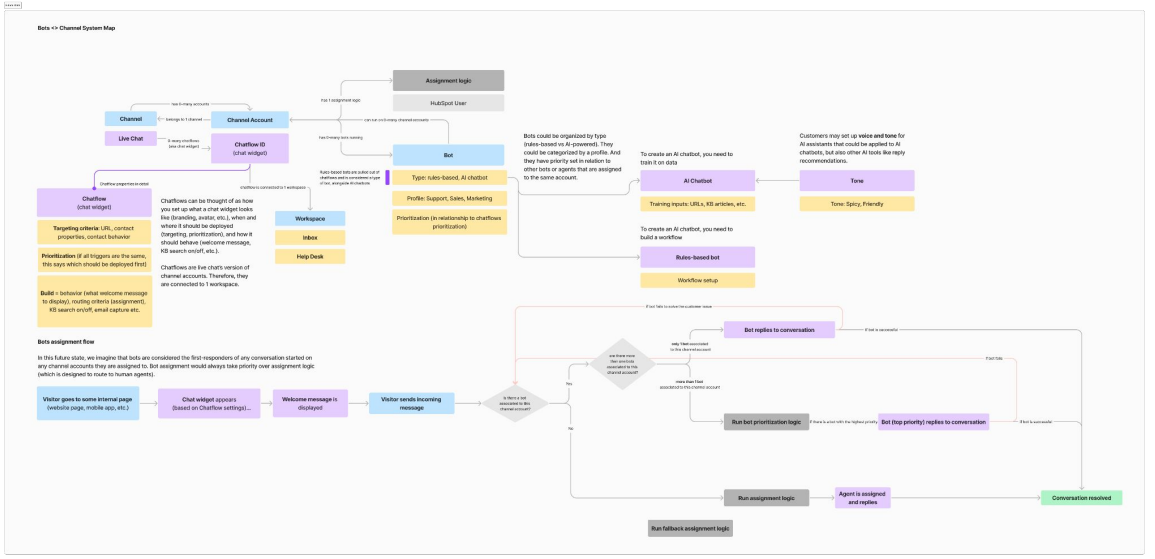

LO-FIS

How can we build trust and offer transparency?

With a solid foundation of research insights and user flows, I translated our strategy into design concepts. I began with **rapid, low-fidelity sketching** to explore a range of solutions and facilitate **early conversations** with product and engineering.

- How can we build trust and transparency with an index page?
- Should we have multiple agents specialized in content for more accuracy?
- What metrics could bring value to our customers. Ultimately these lo-fis helped to scope our V1 approaches vs. future iterations.

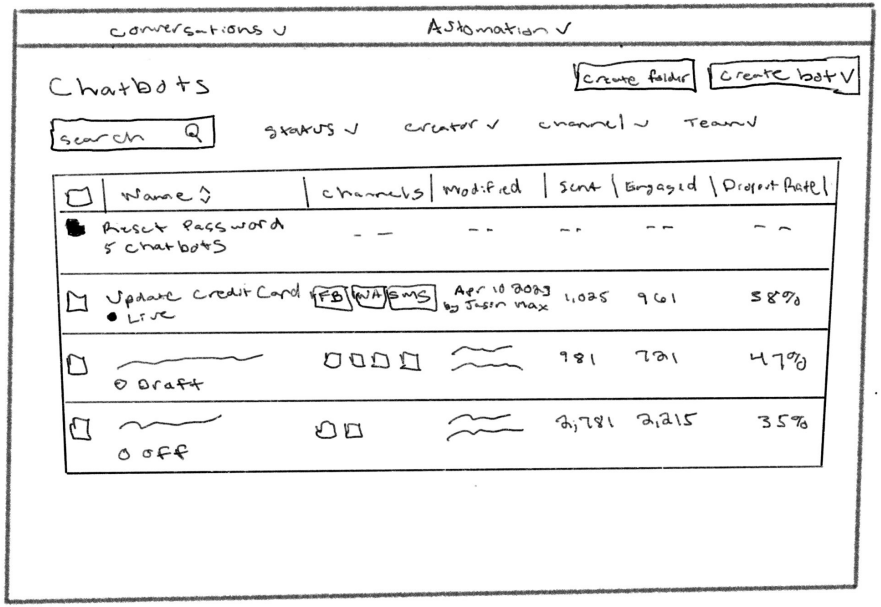

MID-FIS

Customers value trust and are hesitant. How can we educate and seamlessly depoy agents?

After discussions the team came to a consensus that for V1 it would be best to have **1 agent that is intelligently trained to support all potential conversations**.

The next key piece from our research was to **establish trust and provide education** to our customers. I explored potential solutions for what this experience could be?

Do we need a step indicator? How much training is too much? Does this decrease the time to value?

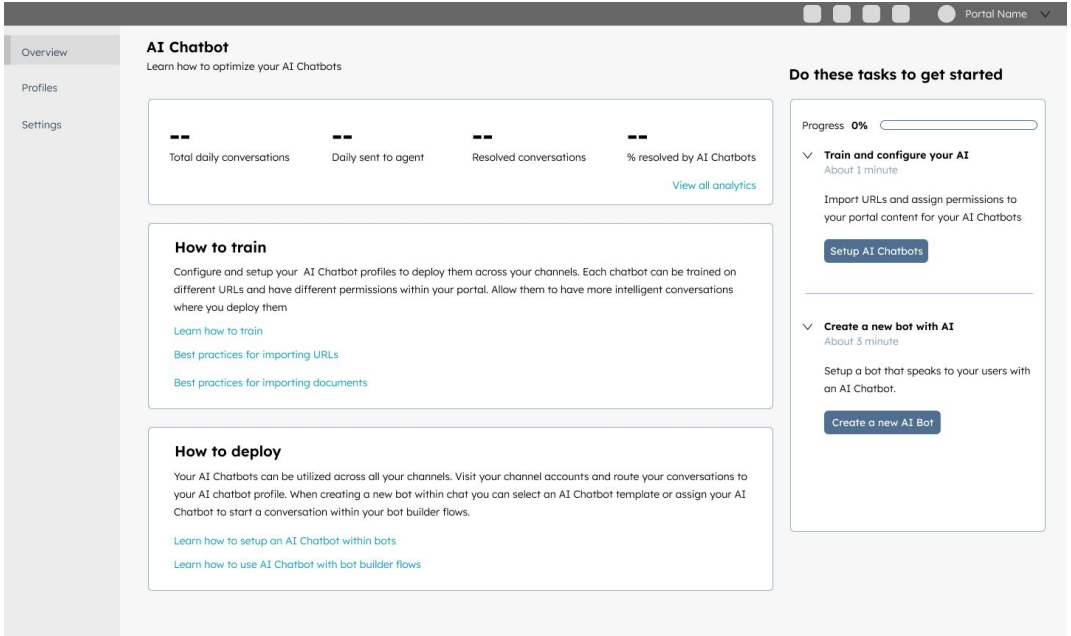

MID-FIS

Transparency: How can customers ensure their AI agent is capable?

Setup, managing, and training were the 3 core pieces to an AI agent. It was crucial that we explored how our customers could seamlessly train their AI.

Should we **auto-train it on Knowledge Base** articles when they are present? We should allow for **customization and flexibility** with their resources. How could an inline training look for quick resource importing?

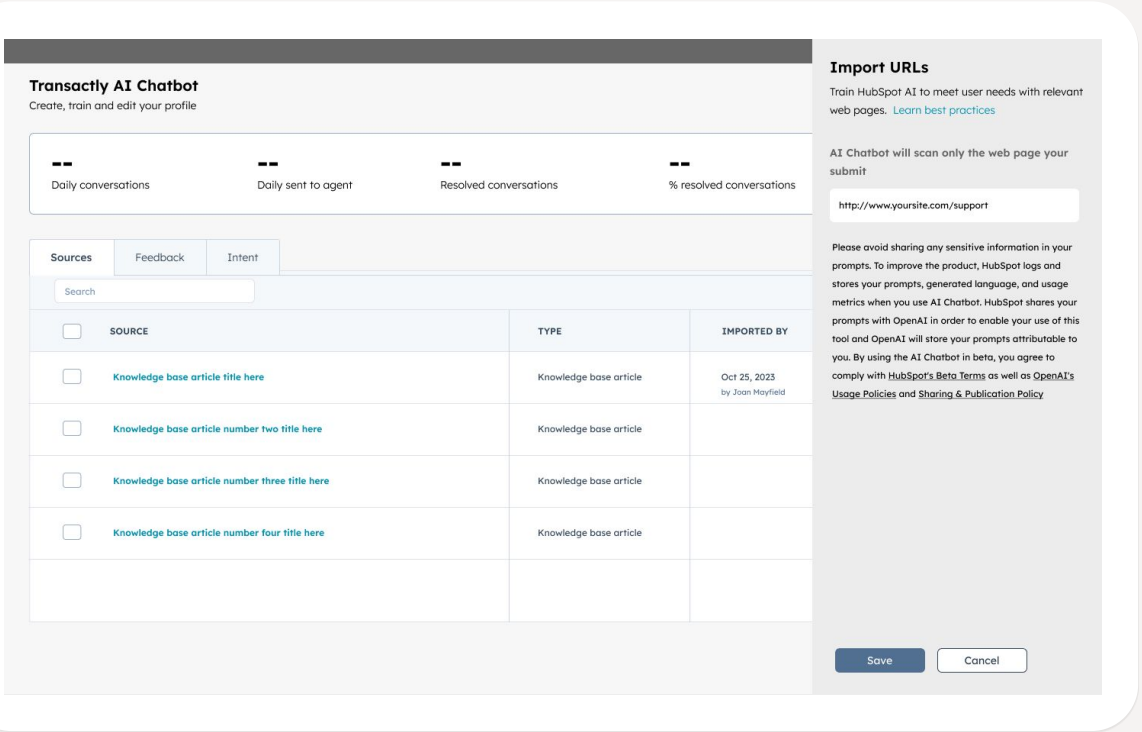

CONCEPT TESTING R1

Are these designs intuitive and effective in offering education and onboarding?

After rapidly exploring a few directions we rescheduled calls with our customers from discovery.

With the **intent to start to validate** our design directions we concept tested two approaches. The focus was to ensure we were properly thinking about **how to educate our users** on this new technology.

Is the training process intuitive for customers? Does it accomplish our **goal for a quick time to value of setup** while offering up education?

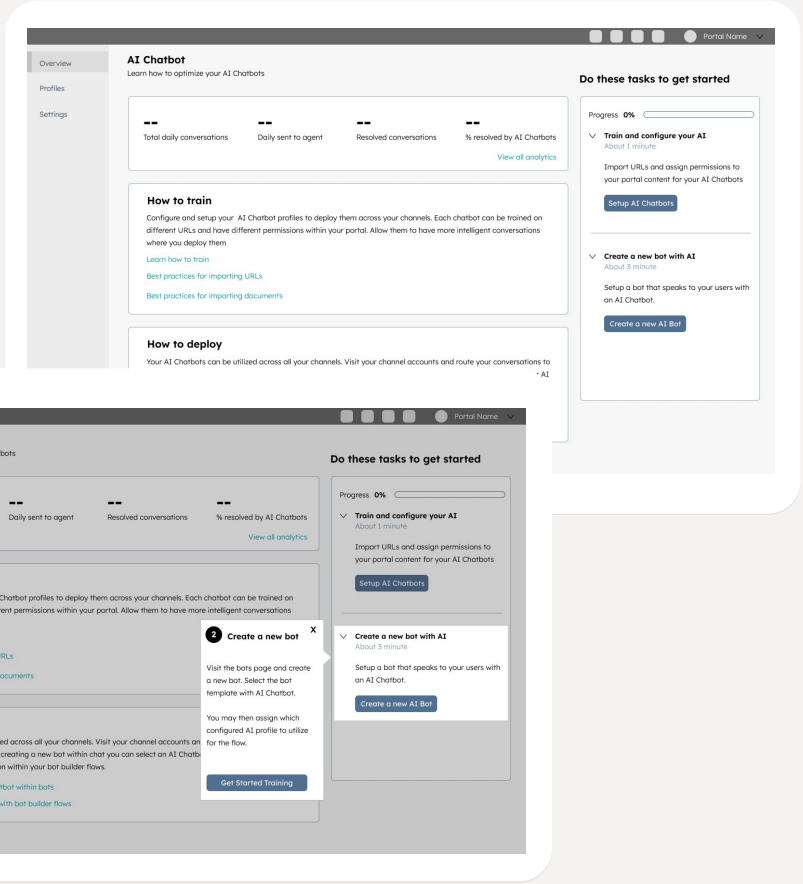

TESTING RESULTS R1

🤔 🤦🏽‍♀️Where did we fail and not meet customers expectations?

Our testing provided some critical feedback on the approach. While customers were encouraged we were exploring how to educate them, **many were overwhelmed**.

What went wrong?

The index page for the AI agent provided **too many resources for education**. Customers were often confused on where to start. They did not know **if they should setup/assign the AI first or train it**.

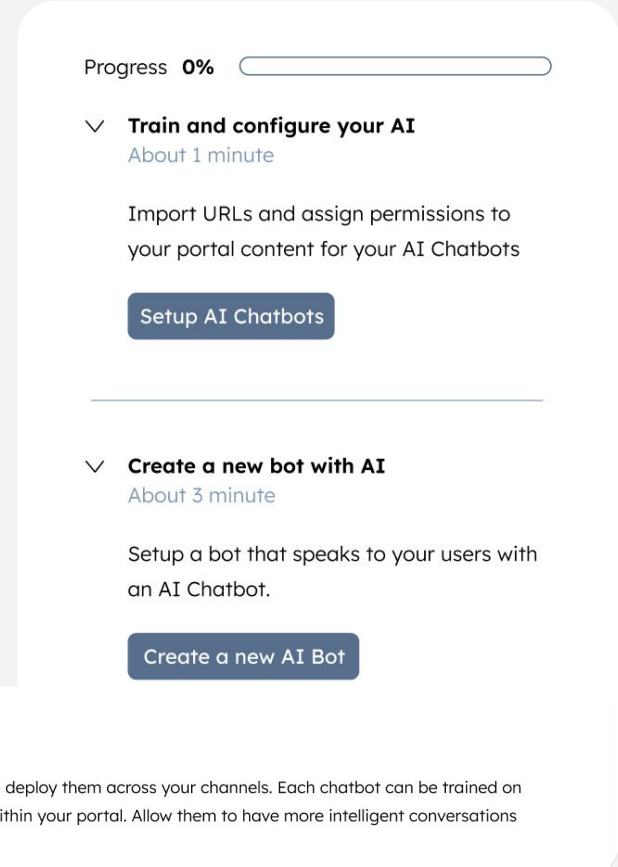

CONCEPT TESTING R2

Based off customer feedback — How do we find a balance between training and education without overwhelming?

Pivoting from the **customer feedback** I explored a new direction and approach to onboarding and education.

It was clear that customers did not need the level of guidance we intended. V2 approach for the mid-fis was to **combine the two steps into a seamless** experience and **reducing cognitive load**.

Allowing customers to gain education while **training their AI and assigning it** to a live chat simultaneously. This would be achieved by integrating the training into the create flow of a live chatflow.

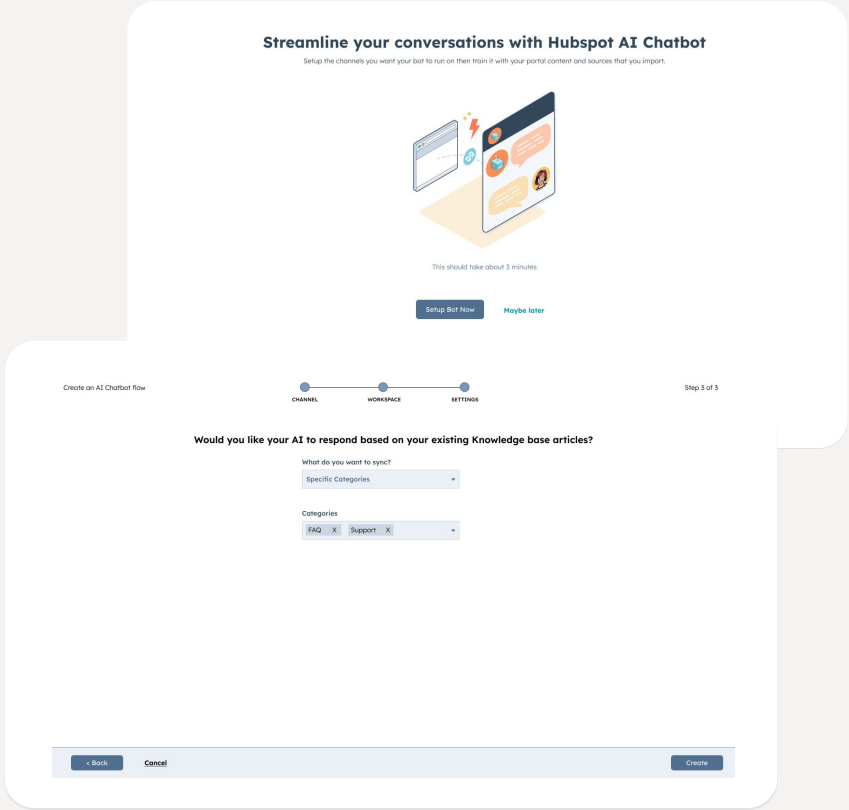

TESTING RESULTS R2

Almost there providing clear onboarding. More robust training and customization is lacking.

Another round of testing offered up additional valuable feedback from our customers. This time around **additional flexibility and customization**.

Customers wanted added level of customization even if that meant additional time spent during setup.

Another key piece of feedback was around onboarding. Our testing revealed that for a new technology like AI, a seamless flow isn't enough; users also **need a clear mental model** *before* they begin.

While the previous version offered too much education, **this version was too simplified**. We needed to find a happy median.

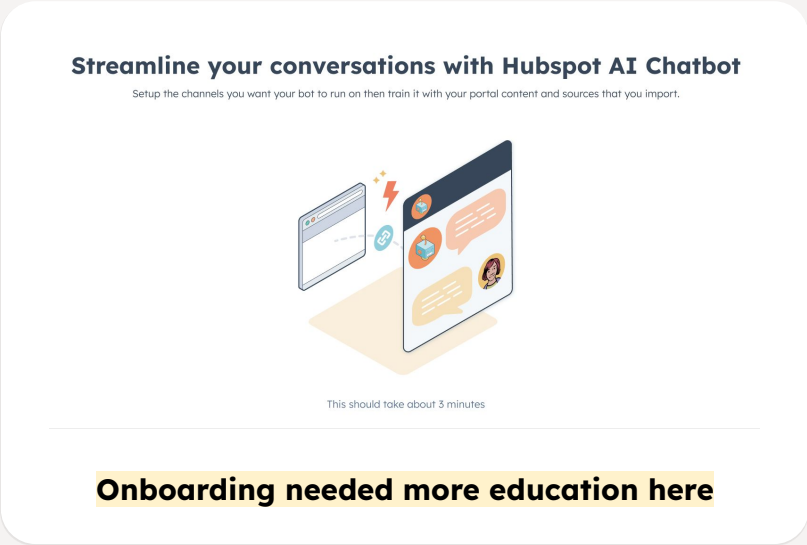

Onboarding needed more education here

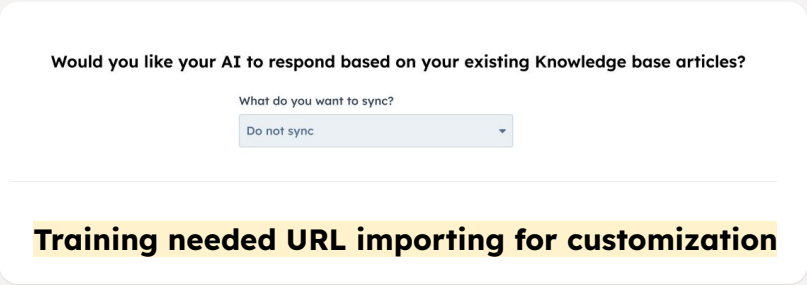

Training needed URL importing for customization

PRIVATE BETA

Time to learn and validate within our system.

Our path to a public launch was driven by the **iterative user feedback**. We began with a private beta, which served as a **live usability lab to test our training capabilities** with customers. The insights gathered during this phase were crucial for refining our backend experience with the temporary FE pages.

Once we felt confident with our BE training capability we continued **validating the value with customers**, then launched to public beta to collect a wider range of feedback and ensure the enhanced AI Chatbot was **ready** for the broader market.

This meant progressing to the **full onboarding and training** flow we gathered feedback on during concept testing.

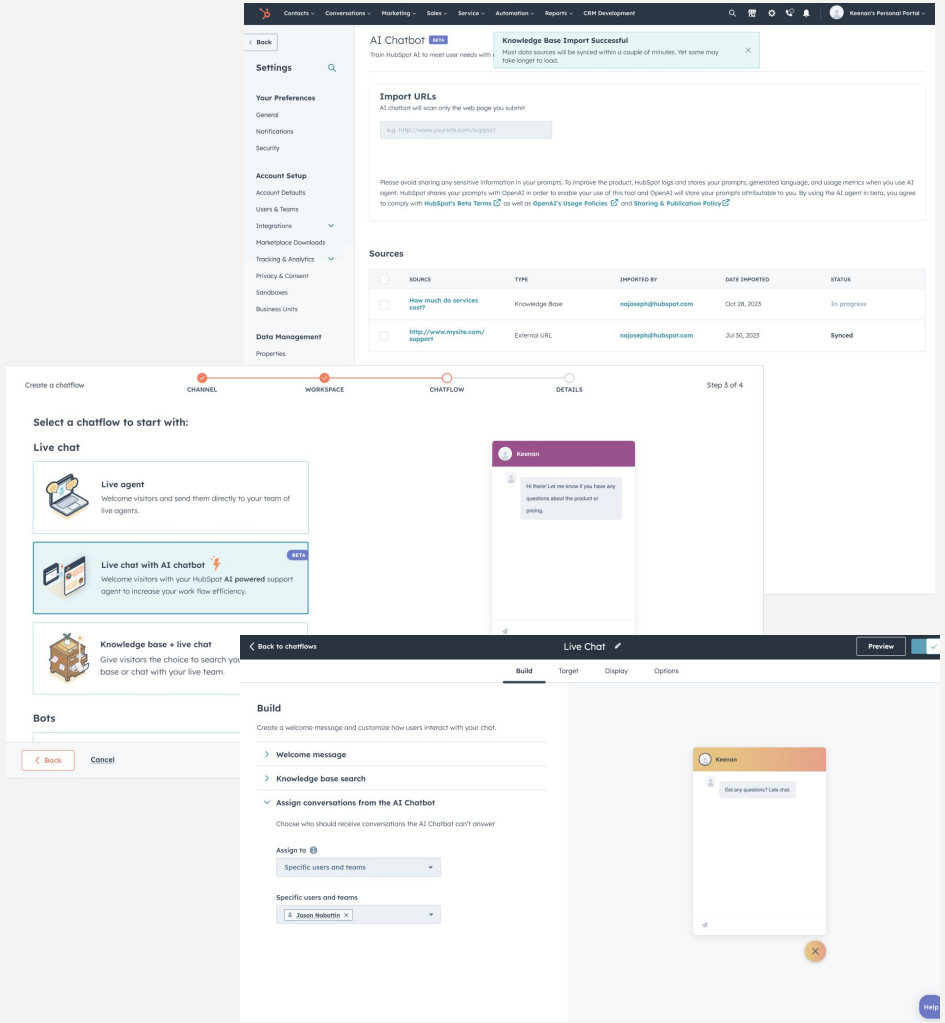

Solution

PUBLIC BETA

Private Beta taught us a seamless setup was crucial for adoption

The primary learning from our private beta and concept testing was that **seamless setups** and **user trust, not just functionality**, was the biggest **barrier to adoption**. Customers repeatedly voiced concerns about AI making things up.

To address this head-on for the public beta, I redesigned the onboarding, setup and management flows to be **transparent and seamless**. We prioritized a **combined setup and training flow** that provided **clear education** on how the AI uses their knowledge sources and built in **safeguards** to give customers confidence in the quality of the responses before going live.

PUBLIC BETA

Building trust & education through onboarding

Recognizing that **trust was one big hurdle**, the first and most critical component of the high-fidelity experience was a strategic landing page.

I designed a dedicated welcome screen that did more than just announce the feature; its purpose was to **proactively frame the user's mental model** for AI.

By clearly **articulating the value**, explaining ***how*** it uses their data to stay grounded in fact, and setting clear expectations, this initial step was designed to shift the user's mindset from **uncertainty to confidence** before they even began the setup process. We identified this as **crucial for driving long-term adoption and retention**.

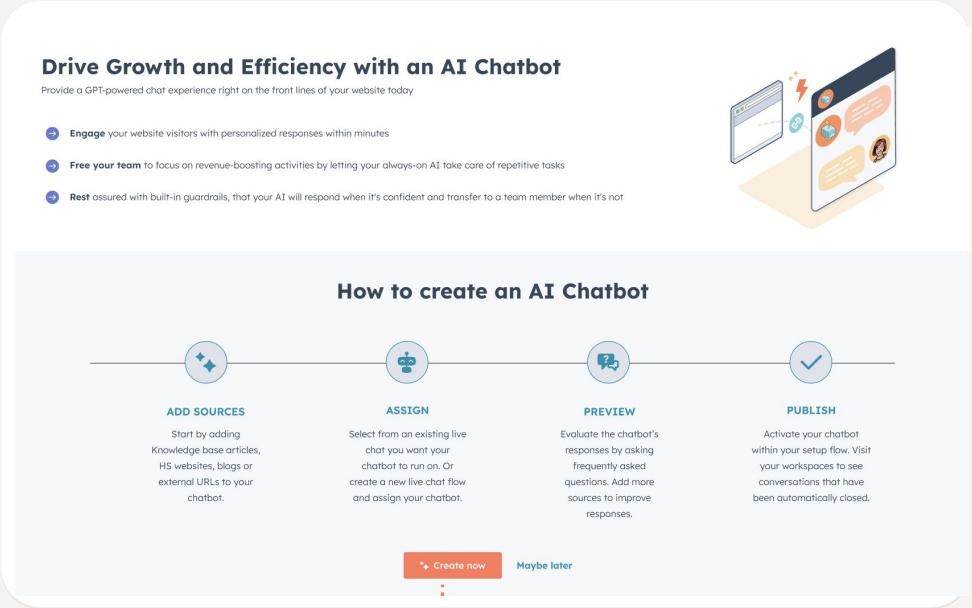

PUBLIC BETA

Providing a seamless setup and flexible customization

After initial education, I designed an intuitive two-step setup flow to get the **AI agent deployed**. The entire workflow was designed to give the user a feeling of control and accomplishment.

Step 1: Train & Customize. This first step **streamlines** the training process. I designed a UI where users could instantly connect knowledge sources via one-click integrations or by simply adding URLs.

Step 2: Assign & Deploy. Once trained, this final step makes the AI actionable. Users can assign the agent to specific chatflows and deploy it with a single click. The design provides **instant visual confirmation** of success, delivering a sense of **accomplishment and a transparent path** to going live.

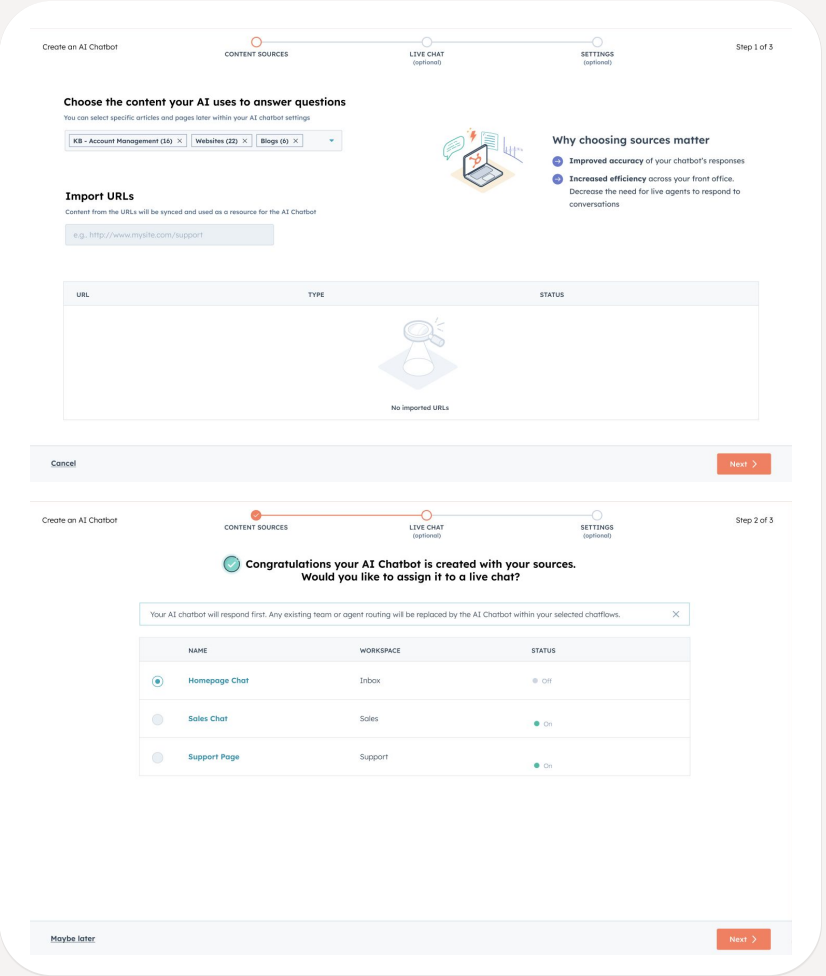

PUBLIC BETA

Trust & transparency through co-piloting an AI agent.

To **build user trust**, we knew the AI sources needed transparency. Based on the principle that a reliable AI requires high-quality data, I redesigned the configuration page to serve as a **transparent command center** for the user's knowledge sources.

This new design gave users **explicit control** over what data the AI could access and, crucially, provided a clear view into the specific sources being used for its answers, **empowering them with both visibility and control**.

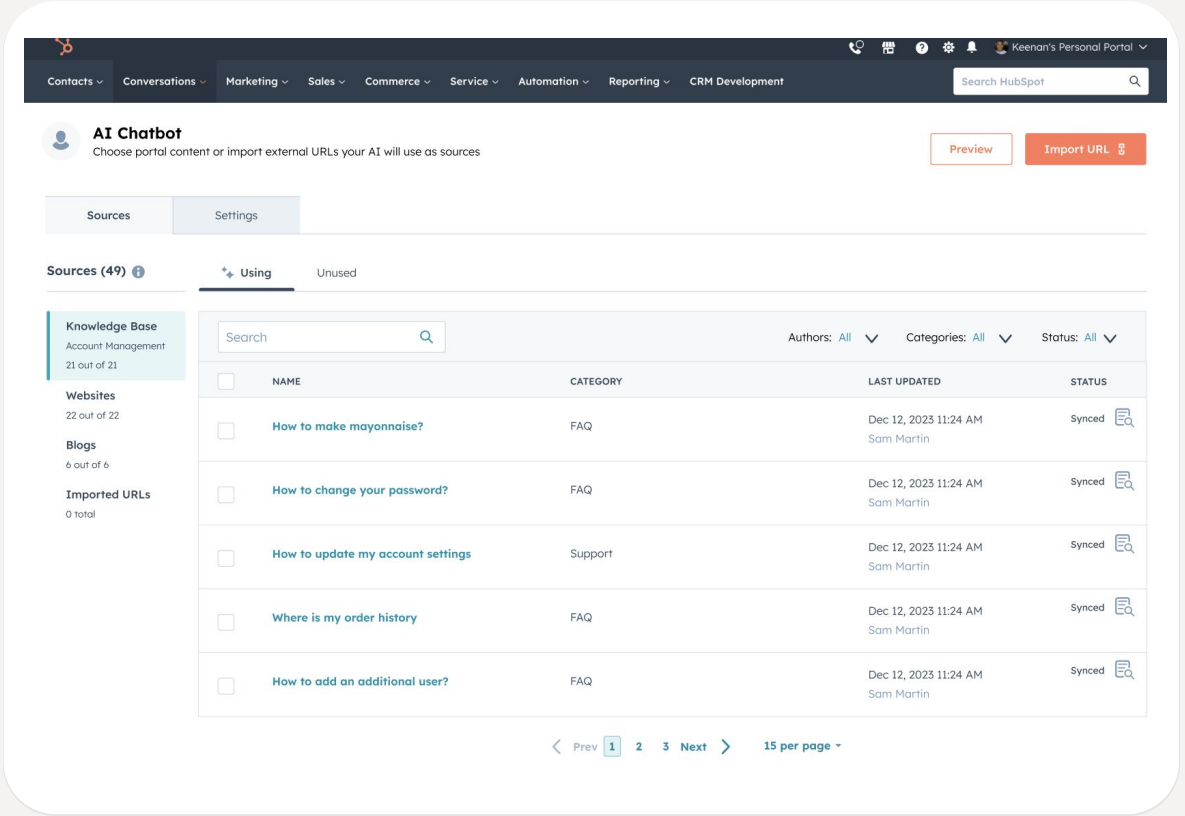

RESULTS

ADOPTION AND RESOLUTIONS

The AI Chatbot demonstrated immediate and significant value following its launch. In the first 60 days, overall adoption of the feature surged by 41% month-over-month. More importantly, customers who participated in our beta program saw their ticket resolution rates jump from an average of 44% to an impressive 64% within their first 60 days of use.

LEARNINGS

KEY LEARNINGS

Continuous Improvement Requires a User-Driven Feedback Loop

A major takeaway was that the AI's launch is just the beginning of its **learning journey**. We discovered that to truly build trust and improve performance, customers need a **robust, in-context system to flag incorrect answers** and reinforce correct ones. For future iterations, I would prioritize designing a **seamless feedback mechanism** that empowers users to **actively co-pilot and train** their AI on the fly.

The Critical Impact of Data Hygiene

We learned that the quality of an AI's output is directly tied to the **structure and quality** of its input. While our system was powerful, its effectiveness was sometimes **limited by poorly organized knowledge bases**. This highlighted a significant opportunity to go beyond the tool itself and **invest in proactive user education**, creating best-practice guides and **in-app suggestions** to help customers structure their data for optimal AI performance.