


How to Deploy AI to Reply to Customer Service Email and Chat



How to Install AI to Reply to Customer Service Email and Chat

- AI has captivated us all—we can't unsee it
- Implementing and governing AI is challenging
- Deploying AI requires flexible workflow orchestration and a nimble integration strategy



Executive Summary



Advances in AI have revolutionized our expectations of human-to-computer interaction. Its sudden surge in popularity caused executives to wonder how to effectively utilize it within their businesses. This guide aims to provide a comprehensive overview of how an AI orchestration platform, specifically Krista, can be used to read intent from customer emails and chats, look up information from your systems and reply with relevant responses to improve the overall

Table of Contents

Customer Service Organizations are Overwhelmed	3
Generative AI Makes it Look Easy	3
Can I just use ChatGPT or Bard in my enterprise?	4
How can I use my company's data with LLMs?	5
Implementing AI Across Systems is Challenging	7
The Intricacies of Enterprise Implementation	8
How to Implement AI to Reply to Customer Emails and Chats	9
Evaluate how customers contact you	9
Define the next-best action based on each category	10
Prioritize categories based on volume and time-to-value	10
A human in the loop builds trust	10
Case Study	11
Revolutionizing eCommerce Customer Support	11
AI-Driven Customer Service Automation	11
The Impact of AI on Customer Service and Operational Costs	11
Krista AI-led Intelligent Automation Platform	12
AI's Impact on Customer Satisfaction and Efficiency	13



Customer Service Organizations are Overwhelmed

Customer service organizations are increasingly overwhelmed with a myriad of challenges. Resource limitations, agent burnout, inefficient workflows, and limited staff pose significant issues. With a constant influx of emails and chats, agents can often become overwhelmed and stressed, leading to burnout and a decrease in service quality. Manual processes for sorting, prioritizing, and routing emails can waste valuable time and lead to delays in response times. Additionally, hiring and retaining qualified agents can be arduous, particularly during peak periods.

Customers' high expectations add another layer of complexity. They demand quick and efficient responses, often within minutes or hours. Failure to meet these expectations can lead to frustration and dissatisfaction. Despite operating in a high-volume environment, customers still crave personalized attention and solutions tailored to their distinctive needs. They also seek consistent communication throughout the resolution process.

Technical challenges further complicate matters. Managing large inboxes and ensuring they do not become overwhelmed can be daunting, especially during peak periods. Protecting sensitive customer information and ensuring compliance with data privacy regulations are paramount. Integration issues, such as combining email with other customer service channels and systems, can be intricate and time-consuming.

AI, specifically platforms like Krista, stands dramatically improve responses by offering solutions to these overwhelming challenges. By reading intent from customer emails, looking up relevant data in your systems, and replying with pertinent responses, AI can significantly enhance the customer service experience.

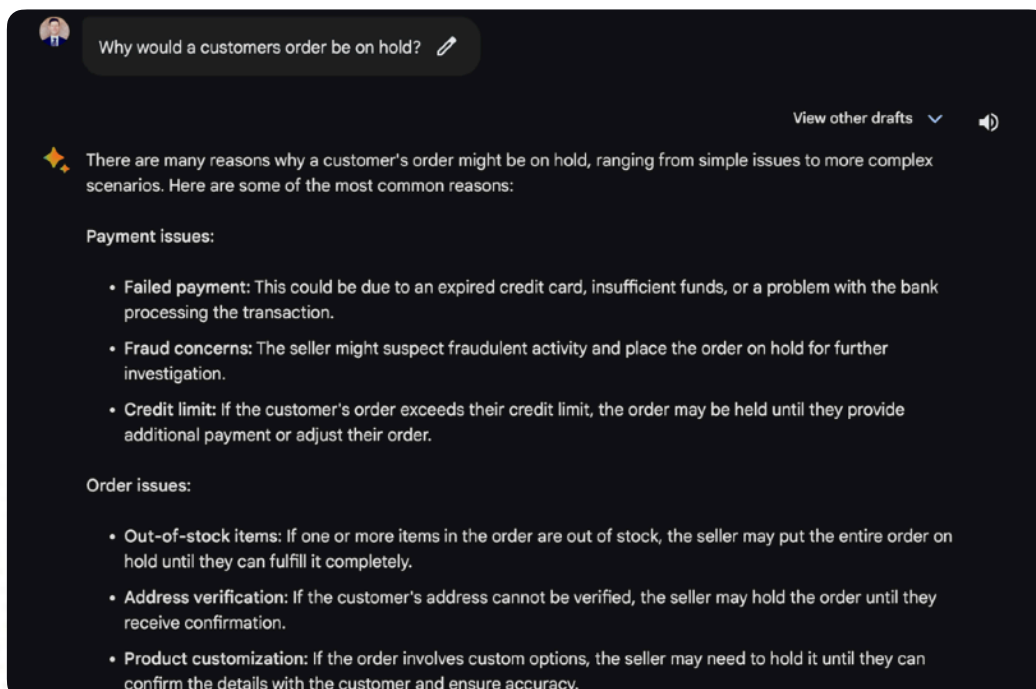
Generative AI Makes it Look Easy

Generative AI (GenAI) is a type of artificial intelligence that uses machine learning to produce fresh content based on existing information. Unlike conventional AI which demands explicit rules and instructions, GenAI has the autonomy to learn and create, proving it valuable for customer service interactions. Advances in Natural Language Processing (NLP) and computer hardware have supercharged GenAI's capacity to comprehend and respond to human communication using everyday language. As such, OpenAI's ChatGPT and Google's Bard have considerably elevated GenAI's visibility since these models offer limited free usage to a wide

audience ranging from customer service representatives to company board members. Such exposure has led executives to experiment with it, imagine potential uses, and wonder how their company business processes can be improved by removing some of the inherent complexities when trying to solve complicated customer service inquiries. This exposure reverses typical technological innovation cycles often initiated by technologists and R&D researchers to business executives who are pressuring CIOs to find profitable use cases for GenAI. However, integrating GenAI into business operations isn't without challenges. Legacy technology, financial constraints, limited resources, hallucinations, and speed of innovation often impede integrating AI within IT infrastructures and workflows.¹

Can I just use ChatGPT or Bard in my enterprise?

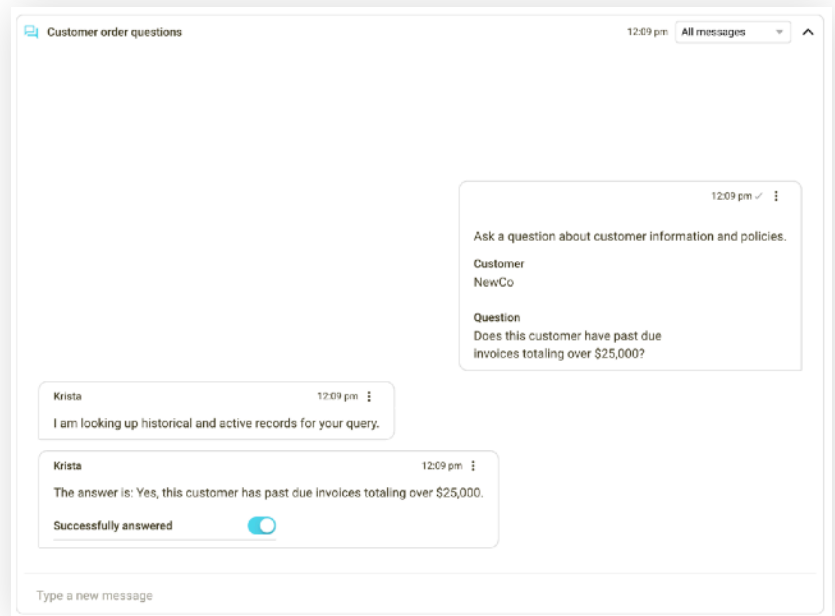
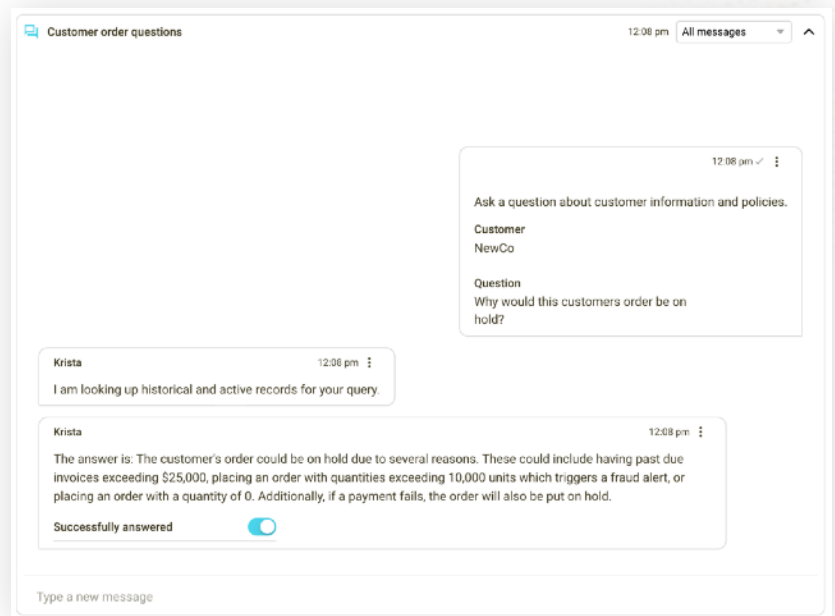
AI solutions for your business must deliver answers based on your company's unique data. Otherwise, public services like ChatGPT and Bard would do just fine but they lack access to your specific data and systems. When posing customer service questions like "Why would a customer's order be on hold?" to ChatGPT or Bard, you receive generic answers derived from their internet-based training or synthetic data, which might be accurate but not tailored to your employees' needs. As you can see from the response screenshot to our agent's query based on a customer service interaction, LLMs like ChatGPT and Bard provide helpful responses but only list possible reasons based on their internet-based training data.



How can I use my company's data with LLMs?

Effective use of AI must combine data from company-specific knowledge, understand the requester's context, and the situational analysis in a prompt. You will need to prompt an LLM with the contextual data so it can generate a response. The LLM will not know the answer. You will not train it on how to answer your agents' questions. For an LLM to generate accurate answers, you will need to supply it with your real-time business data to generate an accurate answer. This is referred to as retrieval-augmented generation (RAG). Retrieval-augmented generation is a technique used to enhance the accuracy and reliability of LLMs by automatically supplying them with data and information from sources like your customer relationship management (CRM) system, knowledge bases, or other internal systems pertaining to a business process to augment responses.

Here is an example using Krista to answer the same question as above but this time retrieving information pertaining to the agents real-time situation.

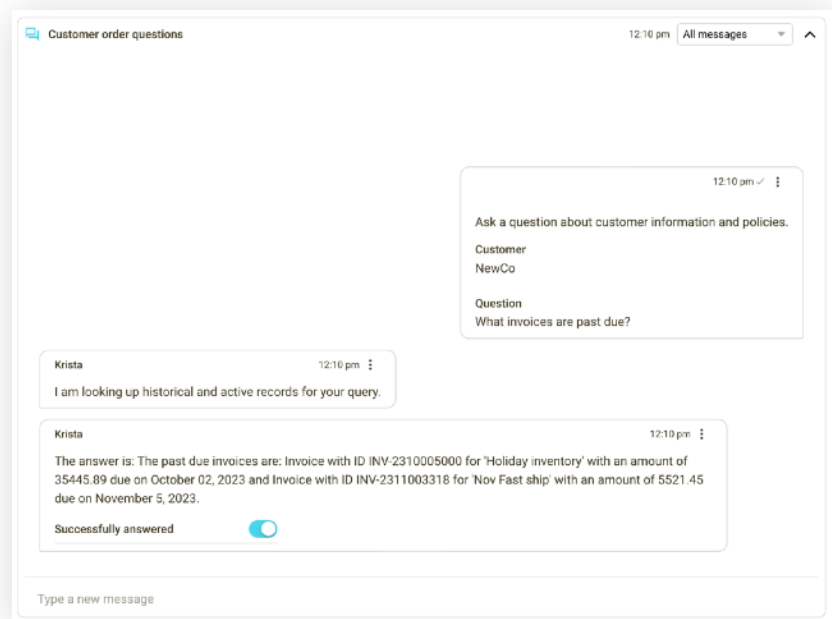


Krista, Krista Software

It starts with a customer attempting to purchase new items. Suddenly, they receive a message informing them that their order is on hold and urging them to contact customer service. The customer reaches out to customer service, explaining the issue they're confronting — their pending orders are on hold. The agent, in response, seeks assistance from Krista, asking, "Why would customer orders go on hold?"

Krista then retrieves information from FAQ and knowledge base articles to provide an answer rooted in the company's actual policies. Krista's response is, "The customer's order could be on hold due to several reasons. These could include having past due invoices exceeding \$25,000, placing an order with quantities exceeding 10,000 units which triggers a fraud alert, or placing an order with a quantity of 0. Additionally, if a payment fails, the order will also be put on hold." To ascertain if overdue amounts are causing the hold, the agent asks, "Does this customer have past due invoices totaling over \$25,000?" Krista, acting as an intelligent assistant, pulls relevant customer data from the CRM and payment systems, reviews invoices, and generates a spot-on, timely response: "The answer is: Yes, this customer has past due invoices totaling over \$25,000." The agent further asks, "What invoices are past due?" Krista retrieves the necessary information from the accounting system and provides the following answer:

"The answer is: The past due invoices are: Invoice with ID INV-2310005000 for 'Holiday inventory' with an amount of 35445.89 due on October 02, 2023 and Invoice with ID INV-2311003318 for 'Nov Fast ship' with an amount of 5521.45 due on November 5, 2023." As you can see this customer has \$40,967.34 due and exceeds the \$25,000 maximum allowed. Therefore, any future orders will not process until they pay past invoices." The agent then confirms the next steps with the customer, ensuring they understand why their orders are not being processed.



Krista, Krista Software

Integrating Krista across several systems not only saves the agent considerable time but also enhances customer satisfaction by delivering a swift and reliable answer. It exemplifies how AI

can streamline workflows and create value in customer service environments. This concept also applies to inbound emails you may receive from customers where Krista can either respond directly or assist an agent with a response.

Retrieval-augmented generation systems present several advantages that contribute to their effectiveness in customer service applications. Firstly, they improve factual consistency by reducing the likelihood of AI models "hallucinating" or generating false information. This is achieved by augmenting the generation process with information retrieved from reliable sources, thereby enhancing the credibility and dependability of AI responses. Secondly, RAG provides provenance for claims, supplying users with access to the information sources utilized by the AI. This transparency enables users to understand the reasoning behind the AI's conclusions, fostering trust in the system's capabilities. Lastly, RAG systems can be fine-tuned for specific tasks by employing varied knowledge bases and retrieval algorithms, allowing for a customized approach to different types of customer queries. By adapting the AI's resource pool and retrieval strategies, businesses can ensure accurate, relevant, and timely responses tailored to the unique needs of each customer interaction.

Implementing AI Across Systems is Challenging

Natural language processing and generative AI technologies have captivated us with their ability to help generate content, understand large amounts of raw data quickly, and provide contextually relevant answers. However, integrating, deploying, and governing generative AI technologies with your data is far from easy. Without the right infrastructure and technology in place, generative AI applications are typically relegated to isolated point solutions or small pilots that lack scalability. As a result, many businesses fail to realize the full potential of these advanced AI technologies and remain stuck in inefficient processes. To truly unlock AI capabilities like in our Krista customer service example, enterprises need modern integration methods that can quickly integrate and govern any AI with existing systems.

Deploying AI inside your enterprise quickly is a challenging endeavor. Traditional software development cycles are slow and time-consuming stifling the speed of innovation. Most IT teams are near or at capacity and can't fulfill the sudden demand to deploy the different AI tools that the business is demanding. If there is capacity, by the time you select and deploy an AI model or service, the scope will have changed and there will be hundreds of new AI alternatives to choose from. Therefore, you need to be able to quickly deploy and interchange AI services in your workflows and decision points and most enterprise teams have too many legacy systems

and data privacy contingencies to deal with to do this quickly. Additionally, traditional iPaaS solutions are helpful when integrating systems but lack the speed, flexibility, and support for integrating LLMs within the context of data governance, making it difficult to securely integrate and deploy models. It's no wonder that so many organizations are struggling to build generative AI applications into their enterprises quickly and securely. However, with the right platform in place, it is possible to unlock the full potential of AI more quickly and securely.

"These large language models do not solve the problem of disparate data sources. Companies need to address data integration and mastering before attempting to access data with generative AI."⁴

Jeff McMillan, Analytics and Data Officer, Morgan Stanley

The Intricacies of Enterprise Implementation

Deploying AI in an enterprise setting is more than meets the eye. While individuals might find generative AI to be a convenient solution for isolated tasks, integrating AI within a business's process across different departments and systems, demands addressing a series of complex challenges. Using a GenAI public interface to help create isolated content or to answer a question makes it seem like AI is solving 100% of the problem. However, it only covers about 5% of the challenges in a business context when trying to assist agents in helping to answer a customer question or to quickly solve a complex issue. The vast majority of the work comes from:

- **Content ingestion:** Importing data correctly is a massive challenge, especially when dealing with varied content like text, tables, images, and metadata. Properly importing, categorizing, and managing this data is a colossal task that requires precision to ensure you prompt an AI model with the right context and information.

- **Real-time access:** Unlike personal use scenarios, where static data is sufficient, enterprises operate in dynamic environments and require real-time data, which means integrating AI models with many existing systems in a nimble and adaptable method capable of retrieving real-time data.
- **Data security:** Enterprises deal with vast amounts of sensitive data, and any AI model or service must operate securely within existing frameworks, ensuring that access is limited to only the appropriate roles and parties and effectively governs your data.
- **Scalability and cost:** Experimenting with public interfaces is free or inexpensive but deploying these models at scale can be extremely costly so enterprises need to be able to manage these costs and be able to interchange AI services without lengthy IT development cycles.

How to Implement AI to Reply to Customer Emails and Chats

Implementing automation and AI to determine intent and reply to customer emails and chats is one of the fastest time-to-value use cases for AI. However, deploying it requires a strategic approach to produce the most accurate responses. Krista provides a simple and efficient way to deploy AI for replying to emails and chats delivering personalized responses to your customers and reducing the workload on human support agents. Here's how you can implement AI to reply to customer emails and chats with Krista:

Evaluate how customers contact you

The first step in leveraging AI to effectively respond to your incoming emails and chats involves understanding your customer interaction channels. You need to evaluate how your customers contact you - is it primarily via email or chat? Doing so helps you work on the most valuable channel to increase your project ROI. Once you choose a primary channel to start, assess whether you currently categorize customer requests into certain categories. If you do, rank the categories in descending order to understand the number of requests per category. If you do not currently categorize requests, Krista can learn from future interactions and categorize inquiries as they happen, thereby streamlining your customer support process. This thorough evaluation forms the foundation upon which AI is integrated into your communication channels.

Define the next-best action based on each category

After identifying each inquiry category like orders, FAQs, invoices, and product support, you will define the 'Next Best Action'. This involves setting up a series of automated steps that are most appropriate for each type of inquiry. For instance, if a customer has a question about an order, the next best action might be to automatically look up information in the order management or customer relationship management software and reply accordingly. If it's a product support or FAQ, the action could be to generate an answer based on a product manual or a knowledge base thereby replying with the information your customer needs. For product support, the action could be to send helpful information tailored to your customer and generate a support ticket to log their request for any warranty or follow-up actions. After defining these actions, they should be prioritized based on the number of inquiries received in each category and the potential value of automating responses. This ensures that the most common and impactful inquiries are addressed in an efficient, timely manner.

Prioritize categories based on volume and time-to-value

Prioritizing categories based on volume is an effective strategy to figure out where to start with the automation process. Categories that generate a high volume of inquiries should be given precedence as their full automation could significantly both increase efficiency and reduce the workload on support agents. For instance, frequently asked questions or common order inquiries might be high-volume categories that would be more valuable to automate. Once you've identified these priority categories, the next step is to determine which systems contain the relevant data or answers about each category. These systems could be your CRM, order management system, or knowledge base. Connecting these systems to the automation process ensures that the AI can access and use the information to generate accurate responses. Once this is done, you can create a tailored conversation to automate each category based on its prioritization and the value it brings to your customer interaction process. This structured approach ensures a smooth, efficient, and value-driven automation of customer inquiries.

A human in the loop builds trust

For those who are concerned with the effectiveness of the AI-driven responses, you can implement a "human in the loop" step in the automation when you get started. This involves creating a step or a gate where a human agent reviews and verifies the AI-generated responses before they are sent. This step helps build trust in the AI system so that all of the stakeholders understand how AI will respond. By comparing the AI responses to those a human agent would respond with, you can validate that the system is functioning as intended. This human oversight ensures that the quality of the service is maintained, and allows you to gradually increase the AI's autonomy as it proves its reliability and accuracy, thereby ensuring a seamless integration of AI into your customer interaction process.

Case Study

Revolutionizing eCommerce Customer Support

A North American eCommerce store was struggling with an expensive and inefficient contact center. The company spent millions annually on a large customer service team that was primarily offshore, leading to extended response times often exceeding 24 hours. The competitive landscape compelled the executive team to seek out opportunities for enhancing operations, streamlining their contact center, reducing costs, and expediting customer resolutions to remain competitive. Observing the growing trend and possibilities of AI, they desired to leverage AI to answer customer queries using data from its ERP system.

AI-Driven Customer Service Automation

The eCommerce store implemented Krista to automate responses to customer emails. Krista applies natural language processing to comprehend a customer’s intent, extract data from relevant systems like ERP, and generate an intelligent response suggesting the next best action. This process greatly speeds up the response time while also reducing the workload on support agents. The team was amazed at how quickly Krista learned from each interaction and continuously improved its responses to match customer expectations.

The Impact of AI on Customer Service and Operational Costs

With the integration and operationalization of AI through Krista, the eCommerce store revolutionized its customer service operations. By providing instantaneous customer responses and accessing necessary customer and order information from the ERP, it ensured swift and precise responses. The organization anticipates a substantial reduction in operating costs as more tasks are automated. Concurrently, this AI-aided model aims to elevate the quality of their customer service, achieving superior resolution times. This advanced approach not only optimizes efficiency and cost-effectiveness, but it also boosts customer satisfaction, marking a significant transformation for the eCommerce store’s contact center operations.



"With the onslaught of GenAI enterprises must proceed at pace to seamlessly integrate a broad set of AI capabilities. Krista's AI iPaaS might be the right starting point."⁶

Tom Reuner, HFS Research

Krista AI-led Intelligent Automation Platform

Krista is a revolutionary AI integration platform designed to easily bring any AI into your enterprise. Krista enables you to easily integrate any AI into your systems and processes to help your people get more done. Krista utilizes a low-code suite of tools enabling you to integrate AI into processes spanning your people, systems of record, data stores, messaging systems, and omnichannel. With Krista, implementing AI into your enterprise systems and workflows can be done with less time and effort than ever before via hundreds of available connectors. Krista's NLP-enabled platform provides natural language capabilities to integrate AI into your existing systems without the need for expensive infrastructure or software development resources. Krista is the perfect fit for unlocking the potential of AI in your enterprise.

The benefits of using Krista include:

- **Scalability:** As a cloud-based service, Krista scales to accommodate increased demand or workloads.
- **Cost-effectiveness:** Organizations can access powerful AI and ML resources without having to invest in expensive hardware and infrastructure.
- **Flexibility:** Krista offers a wide range of tools and services, making it suitable for various industries and applications.
- **Simplified management:** Krista manages your AI applications and resources through a centralized platform, making it easier to monitor performance and make adjustments as needed.
- **Faster time-to-value:** By leveraging pre-built models, low-code tools, and integration connectors, organizations can speed up the deployment process and bring their AI solutions to market more quickly with Krista.

AI's Impact on Customer Satisfaction and Efficiency

Using a structured approach, an understanding of your customer inquiries, and a clear definition of the next actions, it is simple to implement AI into your customer communication channels using Krista. By prioritizing categories based on volume and time-to-value, and maintaining the quality of service with a "human in the loop" approach, organizations like yours can deploy AI effectively for responding to customer emails and chats.

If you would like to use AI to optimize your customer service processes and elevate customer satisfaction, consider utilizing Krista to simplify and expedite the integration of AI into your existing systems. With Krista's powerful capabilities and user-friendly platform, you can revolutionize your customer support operations and stay ahead in an increasingly competitive eCommerce landscape. [Contact us today](#) for a custom demonstration and trial of the most advanced AI-led automation platform.

Sources:

1. [CIOs Feel Heat From CEOs on Generative AI](#), The Wall Street Journal
2. Bard, Google
3. ChatGPT-4, OpenAI
4. [CDO Agenda 2024: Navigating Data and Generative AI Frontiers](#), AWS
5. [Gartner Says Cloud Will Be the Centerpiece of New Digital Experiences](#), Gartner
6. [Krista's platform meets enterprises' urgent need to integrate AI at speed and scale](#), HFS Research
7. [Generative AI for Agile Knowledge Management](#), Krista Software

KristaTM 

