AI Compliance FAQ

Which AI models are used by Kustomer?
Kustomer uses various AI models to power its KIQ offerings. For generative AI capabilities, Kustomer leverages OpenAI across its key products, including Customer Assist, Agent Assist, and AI Agents. In addition to generative AI, other models are utilized to enhance Kustomer's AI functionality. Language detection and sentiment analysis rely on AWS Comprehend generic text-classification models, while message translation is powered by AWS Translate.

Intent detection is supported by two types of machine-learning models:

- A generic model, fine-tuned by Kustomer using freely available and its own curated data and a pre-trained RoBERTa model (published by Meta AI)
- Organization-specific models, trained on a customer's data and based on Rasa, an open-source NLU framework

How does Kustomer handle customer data, privacy, and security in AI systems?
Kustomer ensures responsible handling of customer data by sharing only the content of messages with its AI engines. While most messages lack personal data, Kustomer assumes that some personal data may occasionally be included. To address this, Kustomer has established Data Processing Agreements with third-party subprocessors such as OpenAI and AWS. OpenAI adheres to a Zero Data Retention policy, deleting data immediately after processing, while AWS provides an AI opt-out policy to ensure data deletion.

Customer data is *not* used to train global AI models; however, in cases where AI capabilities are tailored for a specific customer, data may be used to refine models exclusively for that customer's environment. This approach maintains strict privacy standards and complies with GDPR, CCPA, and other relevant regulations.

To further protect privacy and security, Kustomer employs encryption for data both in transit and at rest, implements role-based access control, and conducts regular security audits. These measures are validated by SOC 2 Type II certification.

Additional documentation on Kustomer's privacy and security practices is available on the Kustomer Trust Center.

How does Kustomer ensure AI-generated outputs match brand tone?
When creating or editing an AI Agent, Kustomer provides a place for defining the tone, voice and personality with "Instructions" for the frontline AI agent (a.k.a. Supervisor). Customers instruct the AI Agent in a conversational way just as they would with a ChatGPT prompt.

Customer Assist Premium and Agent Assist use "AI Profiles" to create custom prompts that ensure AI-powered responses align with your brand's unique identity. These profiles allow

administrators to customize the AI's behavior with specific instructions, such as using friendly language or avoiding certain keywords. Localization options enable regional nuances in AI-generated text, and administrators can preview text before implementing it.

Additionally, businesses with multiple brands can create tailored AI profiles for each one.

What guardrails are in place to ensure ethical use of AI?

To prevent misuse of AI outputs, Kustomer allows businesses to set specific guardrails within the platform. These guardrails ensure that the AI operates within defined parameters, such as avoiding discussions about competitors or adhering to other business-specific restrictions. This "do no harm" principle helps maintain ethical and appropriate use of AI.

How does Kustomer inform users about AI-enhanced interactions?

For transparency, Kustomer informs users about AI-enhanced interactions differently depending on the feature. For AI Agents on chat, they display "generated by AI" underneath each message generated by AI. For Agent Assist, there is no explicit indication to customers that messages were enhanced by AI, as these are reviewed and edited by agents before being shared. For Customer Assist, messages generated by AI are labeled as "AI Generated" within the chat widget. Administrators can customize this messaging through the Conversational Assistant settings.

What analytics does Kustomer run, and how is this data used?

Kustomer analyzes aggregated and anonymized data to optimize its services and measure performance. Examples of this analysis include monitoring traffic loads, usage metrics, and evaluating the percentage of conversations automated. This anonymized and aggregated data can also be used in marketing materials to demonstrate service efficacy. This aligns with Kustomer's Master Services Agreement (MSA) which explicitly states that anonymized data may be used to improve services without identifying individual customers.

What licensing terms apply to AI models in KIQ offerings?

Kustomer relies on OpenAI and AWS to provide the foundational models for its AI-powered features. There are no third-party flow-down terms for these models, and customers can also leverage Kustomer's native models for functionalities such as Intent Identification. Supplemental terms for KIQ offerings, including privacy safeguards, are available on Kustomer's website.

Can Kustomer provide documentation and security practices?

Kustomer provides detailed documentation for its AI-powered tools, including Customer Assist and Agent Assist. These documents outline the specific models and services used, such as OpenAI for generative AI, AWS Comprehend for sentiment analysis and language detection, and AWS Translate for machine translation. Security documentation includes information on encryption practices, data-retention policies, and the SOC 2 Type II certification process.

Kustomer's Zero Data Retention policy with OpenAI ensures that data is not stored after processing, providing additional assurance of compliance with privacy standards.