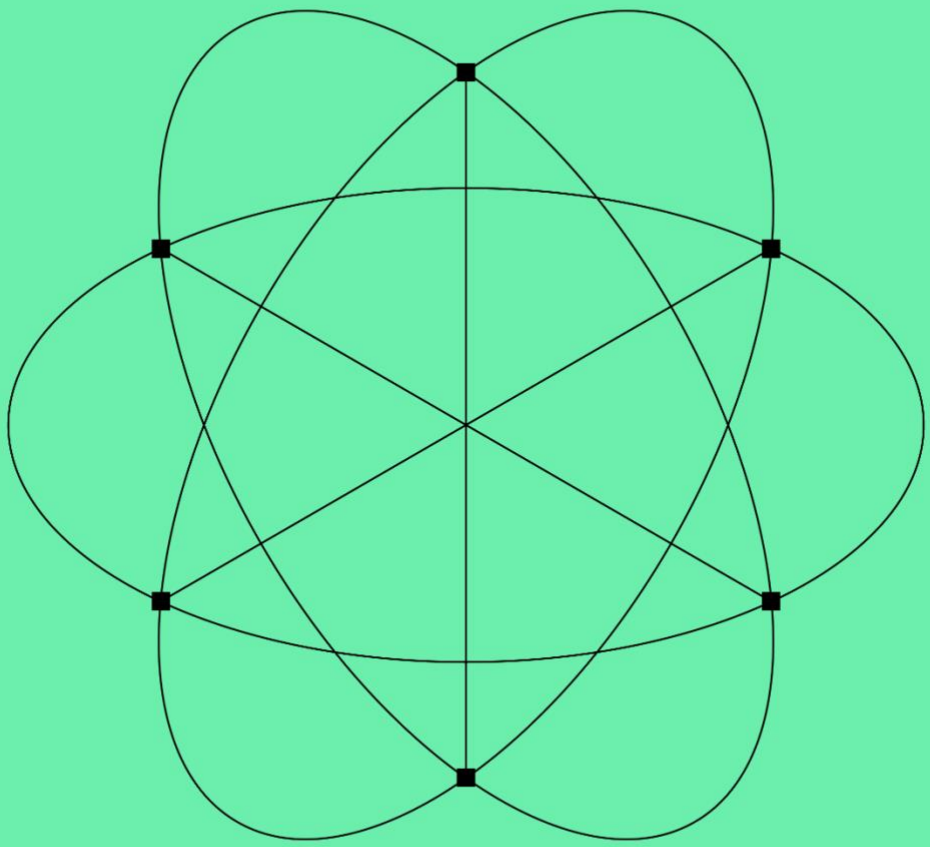
OneTrust

AI Transparency Report

November 2026, V2



DISCLAIMER:

No part of this document may be reproduced in any form without the written permission of the copyright owner.

The contents of this document are subject to revision without notice due to continued progress in methodology, design, and manufacturing. OneTrust LLC shall have no liability for any error or damage of any kind resulting from the use of this document.

OneTrust products, content and materials are for informational purposes only and not for the purpose of providing legal advice. You should contact your attorney to obtain advice with respect to any particular issue.

Table of Contents

Introduction **3**

Description of AI Card **4**

EU AI Act Risk Rating 4

Categorization **5**

A) OneTrust Chatbots 5

AI Feature: myOneTrust Chatbot 5

AI Feature: DataGuidance Copilot (AI Chat Assistant) 6

AI Feature: Platform OneTrust Copilot (AI Chat Assistant) 7

B) Recommendations 8

AI Feature: DSAR Redaction 8

AI Feature: Inventory Record Analysis (Auto-responses) 9

AI Feature: Document Scanner (Auto-responses) 10

AI Feature: Evidence Task Recommendations 11

AI Feature: Evidence Validation 12

AI Feature: Integration Error Analysis 13

C) Content Enrichment 14

AI Feature: Translations 14

AI Feature: Document Classification 15

Introduction

At OneTrust, trust is at the center of our values and that means we are committed to the responsible development and use of artificial intelligence (AI). We are dedicated to ensuring that AI deployed in our application (which we will refer to as AI features) not only meets regulatory requirements but also aligns with industry best practices and OneTrust's AI Principles. OneTrust's AI Principles guide how our teams leverage AI within our solutions:

1. **Transparency by design:** We prioritize openness and explainability through in-product notices where appropriate and updates to our AI Transparency Report.
2. **Privacy-centric data protection:** Transparent data collection, accountable use, and responsible management of data are the key pillars of protecting data in AI features.
3. **Algorithmic accountability:** We require a clear system of accountability for managing AI risks.
4. **Fairness & inclusivity:** We strive to mitigate the risks of unwanted biases and discriminatory practices in our AI features. We believe AI should empower everyone and contribute to the overall prosperity, inclusivity, and growth for all.
5. **Data Stewardship:** Our AI features are designed with industry-recognized information security safeguards.
6. **Safety and reliability assurance:** Our AI features are designed to perform safely and reliably.

OneTrust released our first AI Transparency Report in 2024, and we continue to provide updates as we innovate, learn, and refine our practices. This AI Transparency Report provides information about AI features we have deployed through 'AI feature cards' that outline important information about each AI feature, such as its purpose, methodology, and risk rating. The AI feature cards are organized into categories based on their functionality and data used in development and function.

For more information on our approach to responsible AI see our Trust page.

Description of AI Card

Name of the AI feature, solution, and product: the name of the AI feature is at the top of the AI feature card. The solution(s) and the product(s) the AI feature is associated with are listed in the first row. In some cases, the AI features are applied platform wide.

Purpose: describes the use case and function of the AI feature.

Methodology: provides the name of the AI model used for the AI feature, details on the operation of the AI feature, and information on any training data used in development. The models used are subject to change without notice except when those models process customer data that includes personal data, in which case the process for updating our subprocessors as outlined in our DPA is followed.

Data Processed: outlines the data used in execution to deliver the AI feature. Categories include public and proprietary sources and customer data. Customer data, also known as customer content, means any data or information submitted by or on behalf of a OneTrust customer to the OneTrust platform and applications. Customer data is never used to train the AI models.

Controls: details the options regarding customer control to enable or disable the AI feature.

- **Default:** indicates whether the AI feature is on or off by default when a tenant is provisioned for a new customer (i.e. as the pre-selected option). Default controls are designated by the location of the environment, except dedicated environments, which will not have AI features on by default regardless of hosting location. Where indicated, different default settings may be identified for EU & UK Environments, which include Germany (App-de, App-eu), France (App-fr), Switzerland (App-ch), and United Kingdom (App-uk).

- **Option to Disable:** indicates if the customer administrator can turn the feature on or off.

EU AI Act Risk Rating

As part of our commitment to ethical AI development, OneTrust does not develop ‘Prohibited AI systems’ as defined under the EU AI Act and is not planning to offer any high-risk AI features or systems as part of the services. Below are the descriptions of each AI risk category based on the risk qualification criteria outlined in the EU AI Act:

- **High-risk:** AI features that pose a high probability of harm to health and safety or an adverse impact on the fundamental rights of individuals. Article 6 of the EU AI Act describes the thresholds that lead to an AI system being “high-risk.” Either such a system meets the safety component criteria in Article 6(1) or falls into a category referred to in Annex III (Specific List of High-Risk AI systems).

- **Limited risk:** AI features that use General Purpose AI models to create original content as well as AI features that interact directly with natural persons. Examples include autonomous chatbots; technology that summarizes long-form content, autonomously creates software code, generates digital images from natural language, and produces articles or creative stories based on given prompts; music composition tools that create original scores and songs; and video generation tools that produce animated sequences or edit videos based on text descriptions.

- **Minimal risk:** AI features that do not fall into High or Limited/Transparency AI risk categories. Examples include spam filters, inventory management systems, and business process improvement tools.

Categorization

We organize our AI features into categories based on their functionality and similar data use.

A) OneTrust Chatbots

This category includes AI features that deliver knowledge and answer customer questions. These AI features use natural language interfaces to deliver customized information, oftentimes with links to supporting material included for deeper reference and validation.

AI Feature: myOneTrust Chatbot	
Solution(s) and Product(s)	myOneTrust
Purpose	Conversational AI feature within myOneTrust that helps customers research and obtain tailored answers about OneTrust products, trust information, and related resources, while also supporting customer support interactions and intake. This feature enables customers to access articles, tutorials, support content, and trust information, and helps route or initiate support workflows when customers need assistance.
Methodology	**Which model is used?** Powered by the third party MavenAGI. MavenAGI presently uses OpenAI and Anthropic as model providers. **How does it work?** MavenAGI provides conversational AI and smart search capabilities that process customer input and customer-provided content to deliver chatbot responses and support intake. Customer data is processed for service delivery, and output monitoring is used during deployment. **Training Data** Does not train on customer data.
Data Processed	myOneTrust product documentation and customer-submitted input needed to initiate and support a customer support ticket. Customer input is processed exclusively to provide the requested support interaction and ticket intake.
Controls	Default: On Default for EU & UK & Dedicated Environments: Off Option to Disable: Yes
EU AI Act Risk Rating	Limited risk

AI Feature: DataGuidance Copilot (AI Chat Assistant)	
Solution(s) and Product(s)	DataGuidance
Purpose	Conversational chatbot that answers customers' custom questions about global laws and regulations. Responses are generated from proprietary OneTrust DataGuidance research data. By delivering customized answers to specific customer questions, this AI feature saves customers time on research and finding information across sources.
Methodology	**Which model is used?** Azure OpenAI text-embedding-3-small (third-party), and Azure OpenAI GPT-4 variants (third-party) **How does it work?** A retrieval-augmented generation (RAG) technique converts DataGuidance data into vector representations using Azure OpenAI text-embedding-3-small and then compares the data to find relevant responses to the user's question. **Training Data** Does not train on customer data.
Data Processed	DataGuidance research data. Customer input is processed exclusively to execute the model.
Controls	Default: On Default for EU & UK & Dedicated Environments: Off Option to Disable: Yes
EU AI Act Risk Rating	Limited risk

AI Feature: Platform OneTrust Copilot (AI Chat Assistant)	
Solution(s) and Product(s)	OneTrust Platform (i.e. OneTrust application)
Purpose	A copilot interface for the multiple agentic features that are listed in the methodology section below.
Methodology	**Which model is used?** GPT-4 family from Azure OpenAI (third-party); Gemini-3 flash preview only for Third-Party Risk (third-party) **How does it work?** OneTrust Copilot is an interface or abstraction layer through which the following agentic features can be used: • **Privacy Risk Agent*:** assesses processing activities, suggests assessment templates, and auto-fills assessments; • **Third-Party Risk:** helps complete assessments of third parties; • **DataGuidance:** allows users to learn about and research privacy laws, regulations, and related topics; • **myOneTrust:** answers users' questions about how to use OneTrust products and services; and • **Conversational Analytics*:** answers users' questions about their company's use of the OneTrust platform (e.g., questions about assessments, risks, controls, etc.) and provides tables and visuals in addition to text. **Training Data** Does not train on customer data.
Data Processed	Different types of data are processed depending on the agentic feature used. For instance, for DataGuidance and myOneTrust, myOneTrust product documentation and DataGuidance research data are used. In other cases, for instance, Third-Party Risk, the customer data processed is dependent upon what the customer uploads. For Conversational Analytics, only metadata is provided to the AI models.
Controls	Default: On Default for EU & UK & Dedicated Environments: Off Option to disable: Yes *Agentic features presently in public preview that are denoted with the asterisk above (privacy risk agent and conversational analytics) are default off until they move into general availability.
EU AI Act Risk Rating	Limited risk

B) Recommendations

These AI features provide contextual recommendations to improve the health of compliance at your organization. These AI features never act without user approval and their recommendations are always dismissible.

AI Feature: DSAR Redaction	
Solution(s) and Product(s)	Privacy Automation: Privacy Automation Suite
Purpose	The DSAR Redaction feature detects and allows users to redact confidential information, personal information, and other sensitive information (e.g., financial details and healthcare records).
Methodology	**Which model is used?** Azure NLP v3 (third-party); en_core_web_lg (third-party) **How does it work?** When a document is uploaded, the DSAR Redaction feature converts it into a standard format, and OCR (Tesseract or Azure OCR) is used if the document contains scanned content. Various systems are then initiated to identify the data within the document, and the results are presented to the user in the review step, where the user can validate or edit before approving. Once approved, the DSAR Redaction feature applies the redactions. **Training Data** Does not train on customer data.
Data Processed	Customer data in the form of documents uploaded by the customer
Controls	Default: On Default for EU & UK & Dedicated Environments: Off Option to disable: Yes
EU AI Act Risk Rating	Minimal risk

AI Feature: Inventory Record Analysis (Auto-responses)	
Solution(s) and Product(s)	Privacy Automation: Assessment Automation Third-Party Management: Third Party Risk Management
Purpose	Provides tailored recommendations for assessment questionnaires by analyzing responses from previously completed assessments.
Methodology	**Which model is used?** Azure OpenAI GPT-4 variants (third-party) **How does it work?** Retrieval augmented generation (RAG) is used for this feature. The system first evaluates the list of unanswered questions for an assessment. For each unanswered question, it identifies semantically similar questions from past assessments completed in the same client tenant and incorporates the responses into the final prompt for the large language model (LLM). The system utilizes the LLM's reasoning capabilities to recommend answers for approval by the user. **Training Data** Does not train on customer data.
Data Processed	Customer data in the form of previously completed assessments.
Controls	Default: On Default for EU & UK & Dedicated Environments: Off Option to disable: Yes
EU AI Act Risk Rating	Minimal risk

AI Feature: Document Scanner (Auto-responses)	
Solution(s) and Product(s)	Privacy Automation: Assessment Automation Third-Party Management: Third Party Risk Management
Purpose	Automatically populates assessment questionnaires with suggested responses using information in documents uploaded and scanned by the customer such as a project summary, contract, or SOC 2 report. This saves the customer time and reduces manual effort by auto-filling assessments using insights from added documents.
Methodology	**Which model is used?** Azure OpenAI GPT-4 variants (third-party); Azure OpenAI text-embedding-3-small (third-party) **How does it work?** Retrieval augmented generation (RAG) is used for this functionality. Uploaded documents are turned into vector representations using an Azure AI model and compared to find the ones relevant to the user's assessment question. **Training Data** Does not train on customer data.
Data Processed	Customer data in the form of documents that the customer chooses to upload during an assessment.
Controls	Default: On Default for EU & UK & Dedicated Environments: Off Option to disable: Yes
EU AI Act Risk Rating	Minimal risk

AI Feature: Evidence Task Recommendations	
Solution(s) and Product(s)	Tech Risk & Compliance: all products where custom controls can be created.
Purpose	Simplifies the process of mapping custom controls by using AI to suggest relevant, out-of-the-box evidence tasks.
Methodology	**Which model is used?** text-embedding-3-large (third-party); OpenAI chat model GPT-4o (third-party) **How does it work?** Once a customer creates a custom control library record, this AI feature analyzes the details of the control library record against out-of-the-box controls provided by OneTrust and generates a similarity score to identify the most applicable control. The AI Recommended Evidence Tasks feature then suggests the best evidence tasks to map to the control based on the similarity scores from previous steps. **Training Data** Does not train on customer data.
Data Processed	Customer data in the form of a custom control library record.
Controls	Default: On Default for EU & UK & Dedicated Environments: Off Option to disable: Yes
EU AI Act Risk Rating	Minimal risk

AI Feature: Evidence Validation	
Solution(s) and Product(s)	Tech Risk & Compliance: Compliance Automation
Purpose	Evidence Validation reduces the manual work associated with reviewing evidence for controls by, after scanning attachments, providing a summary of the attachment, assessing whether the contents of the attachment match the requirements of the evidence task (i.e., the relevance and quality of provided evidence), identifying any potential issues or inconsistencies along with remediation guidance, and providing a summary statement on evidence validity.
Methodology	**Which model is used?** Azure OpenAI GPT-4o-mini and Azure Document Intelligence (third-party) **How does it work?** When a customer uploads an evidence attachment (e.g. document) and activates Evidence Validation, this AI feature scans the attachment and generates a list of validation statements or points. The evidence attachment is compared to these validation statements, and a summary of each evidence attachment is generated, including a relevance score (how many of the statements the document satisfies). **Training Data** Does not train on customer data.
Data Processed	Customer data in the form of evidence attachments uploaded by the customer. Evidence attachments may contain personal information, at the discretion of the customer.
Controls	Default: On Default for EU & UK & Dedicated Environments: Off Option to disable: Yes
EU AI Act Risk Rating	Minimal risk

AI Feature: Integration Error Analysis	
Solution(s) and Product(s)	Integrations & Automation within the OneTrust Platform
Purpose	Summarizes failed integration workflow logs by identifying the likely root cause, classifying the issue, and recommending troubleshooting steps, including links to relevant documentation, so users can resolve issues faster and with less reliance on customer support.
Methodology	**Which model is used?** Azure OpenAI GPT-4 family model (third-party) **How does it work?** This AI feature reads integration error logs and uses a GPT-4 family model to analyze the error in context. It then references relevant myOneTrust documentation and generates a structured summary with the likely cause of failure and recommended fixes for the user. **Training Data** None, since it is a pre-trained model. Does not train on customer data.
Data Processed	Customer data in the form of failed workflow log payloads and related event metadata, including request and response details needed to identify the failure, determine root cause, and generate recommended fixes. It may process personal information.
Controls	Default: On Default for EU & UK & Dedicated Environments: Off Option to disable: Yes
EU AI Act Risk Rating	Minimal risk

C) Content Enrichment

This category contains features that enhance or organize content by adding structure, meaning, or accessibility (e.g., translation, tagging, and categorization).

AI Feature: Translations	
Solution(s) and Product(s)	Privacy Automation: Data Subject Requests
Purpose	Reduce time to fulfill requests by automatically translating inquiries sent by data subjects in other languages to the preferred language spoken by your privacy team and enabling data subjects to translate comments with respect to privacy rights requests. In addition, it can translate any comment in a user's session language.
Methodology	**Which model is used?** Microsoft Translator 3.0 (third-party) **How does it work?** A data subject may include a comment when submitting a privacy rights request in the OneTrust Privacy Automation solution. If a comment is not in the viewer's session language, it can be translated using Azure AI Translator, a cloud-based neural machine translation service. The comment is sent to Azure AI Translator through a JSON-based web API, and the translated output is returned to the customer's tenant. This applies both to comments submitted in Privacy Automation and to comments a data subject receives in a customer's Trust Center. The feature can also translate any comment, including internal comments, into the user's session language. For example, in multinational organizations, privacy team members can translate comments posted internally in different languages. **Training Data** Does not train on customer data.
Data Processed	Customer data in the form of data subject comments. Data subject comments may contain personal information.
Controls	Default: On Default for EU & UK & Dedicated Environments: Off Option to disable: Yes
EU AI Act Risk Rating	Minimal risk

AI Feature: Document Classification	
Solution(s) and Product(s)	Data Use Governance: Data Discovery & Classification
Purpose	Categorizes documents into pre-defined types that are common in businesses. As an example, it can determine whether a document is a resume, patent, or contract. Document organization and retrieval is streamlined through this rapid classification process.
Methodology	**Which model is used?** Custom support vector machine (SVM) based model (OneTrust proprietary),Word2Vec or GloVe for embeddings (third-party), and Tesseract for OCR (third-party) **How does it work?** The document type is predicted using supervised machine learning (ML) on text extracted from the file. This process is accomplished by the following steps: the extracted text is first converted to vector representation using text embedding techniques (such as word2vec or GloVe), which are pretrained or fine-tuned; next, based on the text embedding, a SVM classification algorithm is used to decide the document class. Where there are images, Tesseract is used for OCR. **Training Data** Does not train on customer data.
Data Processed	Customer data in the form of documents found during the Data Discovery scanning process.
Controls	Default: On Default for EU & UK & Dedicated Environments: Off Option to disable: Yes
EU AI Act Risk Rating	Minimal risk