



Arctera™ Insight Platform

AI Capabilities Overview

Contents

Purpose:	3
AI Governance and Assurance Framework	3
ISO 42001 Readiness	3
Hallucination Mitigation	3
Toxicity and Content Moderation	3
Bias Mitigation	3
Data Security and Loss Prevention	3
Intellectual Property and Copyright	4
Summary of Arctera AI Capabilities	4
Intelligent Review Machine Learning	5
Model Details:	5
Language Detection in Classification	6
Model Details:	6
Sentiment Detection in Classification	6
Model Details:	6
Voice Transcription	6
Model Details:	6
Text Translation within Surveillance	7
Model Details:	7
Arctera Insight AI Assistant	7
Model Details:	7
InsightBooks	8
Model Details	8
Header: Frequently Asked Questions (FAQ)	9

Purpose:

This document outlines the AI capabilities of the Arctera Insight Platform. All features described are optional and can be configured based on customer requirements.

AI Governance and Assurance Framework

Arctera is committed to transparency, security, and ethical use of AI across its Insight Platform. The following provides context on the certification, risk management, and data governance measures that underpin our AI capabilities.

ISO 42001 Readiness

Arctera is aligning its AI management practices with the ISO/IEC 42001 standard for responsible AI governance. This work includes itemizing compliance controls, documenting lifecycle management procedures, and preparing for future certification.

Hallucination Mitigation

No industry framework fully eliminates hallucinations in generative models. Arctera works to minimize this risk by grounding responses in each customer's archived data. The model retrieves verified content from that archive before generating an answer, aligning facts with actual records. Responses are paired with source links or references for auditability and transparency.

Toxicity and Content Moderation

All InsightAI processing occurs through Microsoft Azure OpenAI within the customer's tenant. Azure OpenAI includes mandatory, built-in content filters (Default V2) that automatically evaluate both text inputs and model outputs.

This built-in moderation system checks prompts and completions for harmful or unsafe content categories defined in Microsoft's Azure AI Foundry Models Content Filtering framework, including:

- Hate or unfairness
- Violence
- Sexual content
- Self-harm

If a violation is detected, the filter either blocks the request or returns a flagged response according to Microsoft policy. These safeguards are always active and cannot be disabled, providing consistent moderation across all InsightAI interactions.

Bias Mitigation

While bias cannot be entirely removed from large language models, Arctera employs mitigation measures to minimize influence on outputs. InsightAI is grounded only in verified customer data, avoiding dependence on public datasets. Prompts and templates are structured using neutral framing and inclusive language to further reduce skew and maintain objectivity in generated text.

Data Security and Loss Prevention

Data used in AI queries remains within Arctera's Microsoft Azure subscription. Query content, system messages, and document context are transmitted securely to Microsoft Cognitive Services for processing and deleted immediately after completion. No customer data is retained, shared externally, or used for model training. These controls align with the principles documented in the Platform Risk Impact Assessment (PRIA).

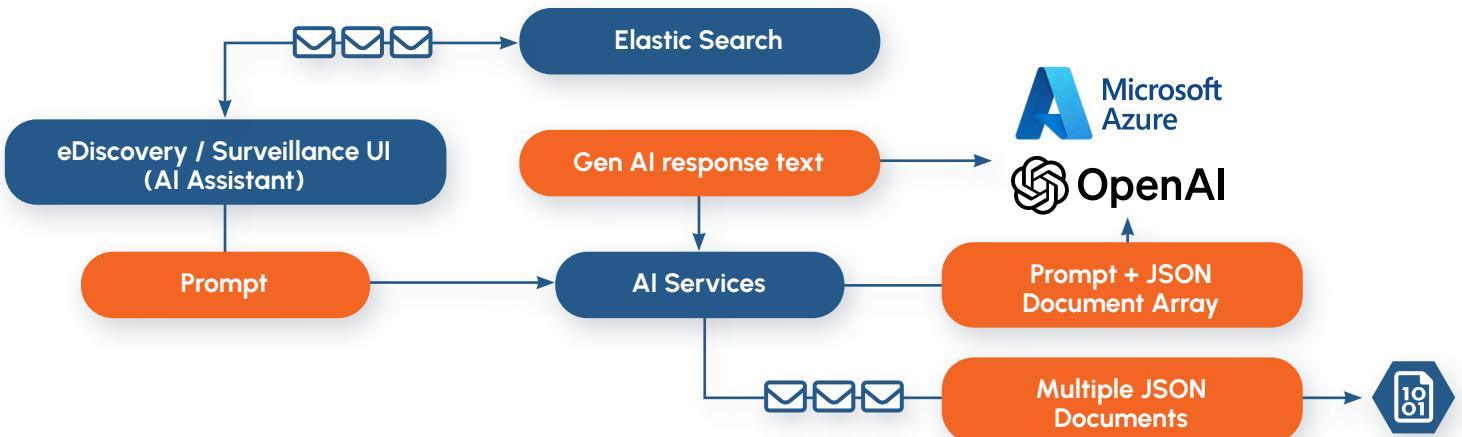


Figure 1: Result Set Level Generative AI

Intellectual Property and Copyright

Customers' content is used only to generate contextual responses, not to train or enhance the model. Generated text does not create derivative works, and no customer material persists beyond its original form in the archive. This always keeps any applicable intellectual property under customer control.

Summary of Arctera AI Capabilities

Capability	Description	Data Stays within Tenant	Model trained by Data Set	Partner Solution?
Intelligent Review	Uses AI-ML to classify items based on previous reviewer actions, helping to determine the relevancy or irrelevancy of content for future reference.	Yes	Customer-specific, Data Encrypted (Department Level)	No
Language Detection (AI/NLP)	Detects the language of text to enable accurate processing of multilingual content.	N/A	Pre-trained model offered by Arctera (Not influenced by customer data)	No
Sentiment Analysis (AI/NLP)	Identifies the emotional tone and sentiment of text (e.g., positive, negative, neutral) for deeper content insights.	N/A	Pre-trained model offered by Arctera (Not influenced by customer data)	No
Voice Transcription	Converts A/V content into searchable text and provides an immersive media player for streamlined review.	Stays within Arctera Microsoft Subscription but external to Arctera DC	Pre-trained model supplied by Microsoft	Microsoft
Translation of Text Content	Converts text-based items, including transcribed content, into selected languages with a single click, like Outlook's feature.	Stays within Arctera Microsoft Subscription but external to Arctera DC	Pre-trained model supplied by Microsoft	Microsoft
AI Assistant (GenAI Queries):	Queries data, answers questions, and provides document and case summarizations, functioning like a customized ChatGPT trained on your data.	Stays within Arctera Microsoft Subscription but external to Arctera DC	Pre-trained model supplied by Microsoft	Microsoft
InsightBooks	Surfaces detailed insights from your personal content archive, organized by topics or subjects, using GenAI trained on your data.	Stays within Arctera Microsoft Subscription but external to Arctera DC	Pre-trained model supplied by Microsoft	No

Intelligent Review Machine Learning

Intelligent Review is built on a continuous active learning framework that updates itself daily. This automated background task observes your review team as they mark and label items for relevance. Since it utilizes reviewer history, there is no need to create a special training set or tune the engine. Essentially, machine learning simply observes your team and learns continuously. You won't need a subject matter expert or a data scientist; the framework is integrated into the system and can be activated with a simple button click.

Unlike most legacy systems that require separate training events and dedicated teams, this system never becomes outdated or drifts toward inaccuracy or irrelevance. Instead, its accuracy improves over time—the more data you provide, the better it becomes.

Model Details:

- Version: Aligned with the quarterly release of the Arctera Insight Platform.
- Developed exclusively by Arctera.
- Based on actions taken by company staff during their review of communications in Arctera™ Insight Surveillance.
- Training affects only your company's data set, currently limited to the Department level.
- There is no sharing of data for external AI model training.
- Operational Boundaries: All activities are conducted within your company's tenant.

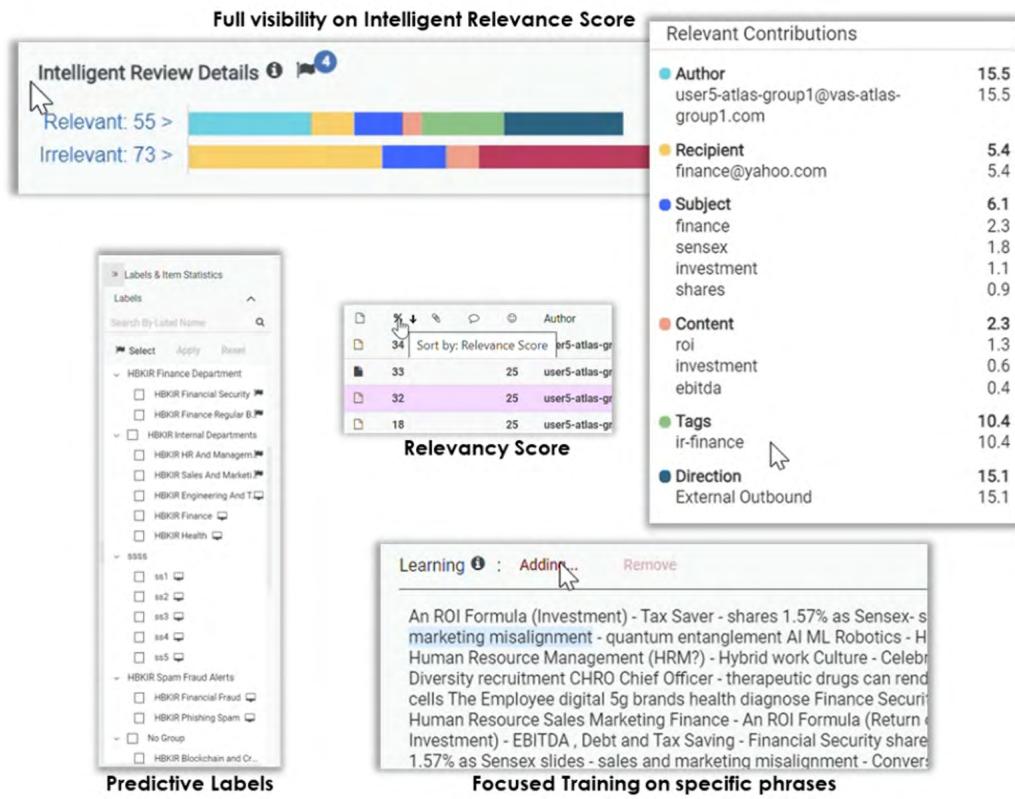


Figure 2: Arctera Insight Intelligent Review machine learning.

Language Detection in Classification

Language detection can detect over 30 languages via AI/NLP (natural language processing). We have specific classification policies for each language and can filter based on languages during review. We also provide a policy that can detect multiple languages in a single message to help detect when someone leverages “Google Translate” or similar technologies that might otherwise bypass detection.

Model Details:

- Arctera developed technology
- Pre-trained model
- No sharing of data for external AI model training
- All activities occur within the Arctera Data Center

Sentiment Detection in Classification

This element of the classification engine is designed to identify the sentiment of the text. The Arctera Classification Engine typically outputs a sentiment score between 0 (Negative) and 1 (Positive). This can be used to guide reviewers to messages with extreme sentiment.

Model Details:

- Arctera developed technology.
- Pre-trained model.
- There is no sharing of data for external AI model training.
- All activities occur within the Arctera Data Center.

Voice Transcription

Arctera Transcription Services offers a comprehensive solution to address compliance challenges in monitoring electronic communications. Our platform provides accurate and efficient transcription of audio and video media files, enabling institutions to safeguard against regulatory threats while enhancing operational efficiency. Arctera Transcription evolves as language evolves by leveraging Microsoft Azure OpenAI to stay current with new world leaders, businesses, and new words entering the lexicon (ex., Brexit, Covid).

Model Details:

- Microsoft OpenAI technology.
- Pre-trained model.
- There is no sharing of data for external AI model training.
- Messages are sent to Microsoft Cognitive Services (Azure Voice OpenAI) for processing; however, the message stays within the Arctera subscription. Once processed, the voice message is deleted, and the transcript is returned.

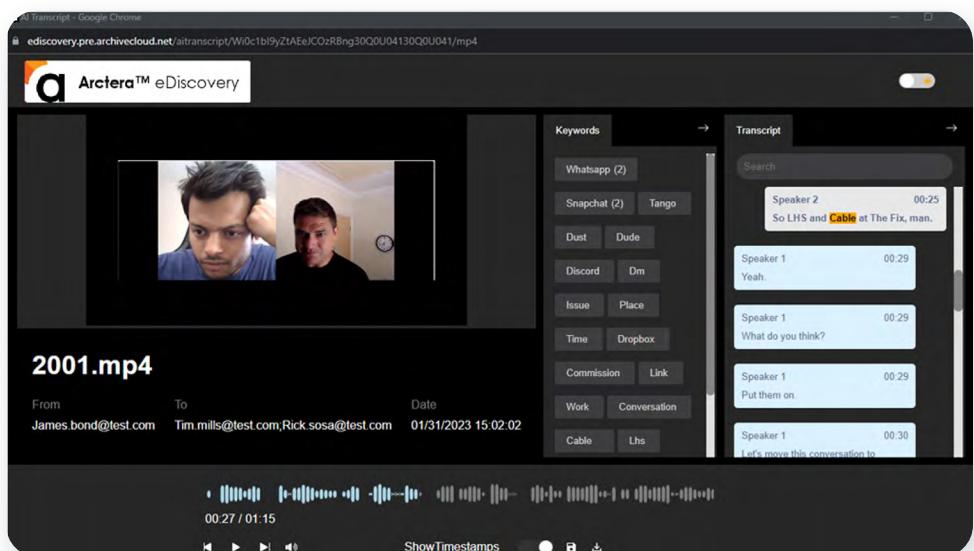


Figure 3: Arctera Insight Native Audio/Video Viewer.

Text Translation within Surveillance

This feature allows conversion of text-based items into a selected base language to allow review within the Surveillance interface. This feature would, for example, allow an English-speaking reviewer to review non-English languages.

Model Details:

- Microsoft Cognitive Services
- Pre-trained model
- There is no sharing of data for external AI model training
- Message body is sent to Microsoft Cognitive Services for processing; however, the message stays within the Arctera subscription. Once processed, the original messages are deleted, and the translated text is returned.

Figure 4: Arctera Surveillance Text Translation View.

Arctera Insight AI Assistant

The Arctera Insight AI Assistant is used to query data sets or even a specific message to gain quick insights. The AI Assistant allows for summarization across a wide data set, topic mining and even summarization by user. Data sets can be interrogated in a similar fashion to Co-Pilot to gain quick insights.

Model Details:

- Microsoft Cognitive Services
- Pre-trained model
- There is no sharing of data for external AI model training
- Query, System Message, and Document context are all sent to Microsoft Cognitive Services for processing. However, the message stays within the Arctera subscription. Once processed, everything is deleted and the response to the query is returned.

Figure 5: Arctera Insight AI Assistant.

InsightBooks

InsightBooks is a capability within Personal Archive that uses GenAI to surface insights from an individual's communication history, including email, Teams chats, Slack messages, and meeting transcripts. It organizes this content by topic, creating a structured view of conversations and decisions over time.

As new content is added to the archive, InsightBooks automatically updates, connecting related messages and highlighting patterns that are often lost in day-to-day workflows. It helps users quickly find relevant context, reduce manual search, and retain visibility into long-running workstreams – all within a secure, governed environment.

Model Details

- Microsoft Cognitive Services
- Pre-trained model
- There is no sharing of data for external AI model training
- Query, system message, and document content are sent to Microsoft Cognitive Services for processing. The message stays within the Arctera Microsoft subscription. Once processed, the original messages are deleted and only the generated insights are returned.

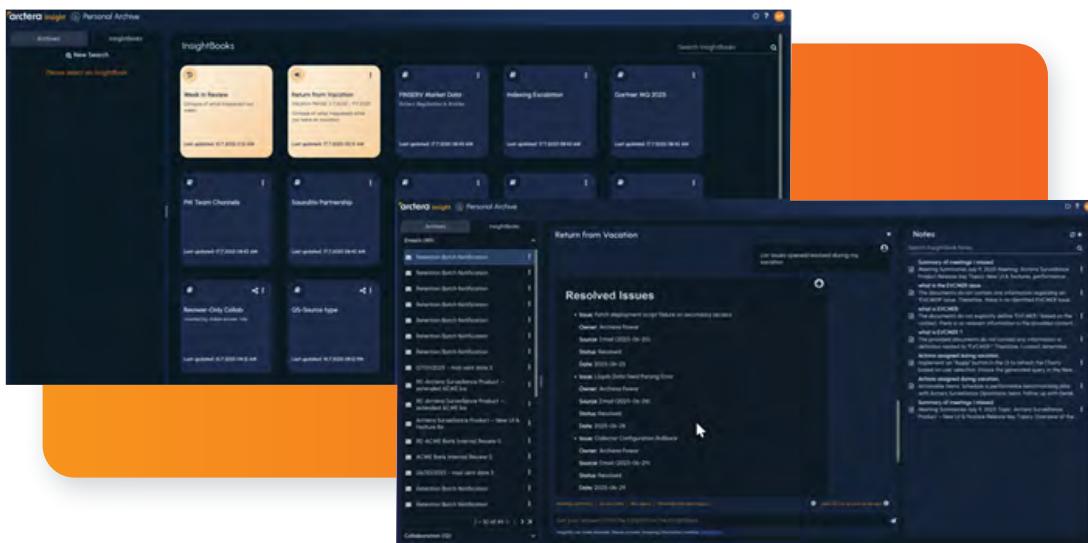


Figure 6: Arctera InsightBooks view, with user communications organized into automatically generated topics.

Header: Frequently Asked Questions (FAQ)

For summary purposes. See the AI Governance and Assurance Framework section for detailed information.

Is Arctera's AI stack ISO 42001 certified?

Arctera is aligning its AI management practices with the ISO/IEC 42001 standard for responsible AI governance. Certification readiness activities are in progress.

How does InsightAI handle hallucinations?

InsightAI grounds each answer in the customer's archived data. Responses are generated from retrieved source content and include reference links where available.

What safeguards prevent harmful or inappropriate responses?

Azure OpenAI includes mandatory Default V2 content filters that automatically review prompts and model outputs for hate, violence, sexual content, and self-harm. Flagged content is blocked or returned according to Microsoft policy.

Does Arctera share my data with Microsoft or use it for training?

No. Query content and document context are processed within the customer's Azure subscription and deleted immediately after completion. No data is used for external training.

Can bias affect InsightAI responses?

InsightAI uses verified internal data and neutral prompting structures to help minimize skew or bias in generated output.

What about copyright or IP exposure?

InsightAI does not retain or train on customer content. Uploaded materials remain within the customer's environment and are used only for generating contextual responses.

About Arctera

Arctera, a business unit of Cloud Software Group, is the leading global provider of compliance and governance solutions that enable firms to unleash game-changing technologies into their organizations while minimizing risk. Created in 2024 from Veritas Technologies, Arctera helps the biggest companies in the world monitor and control exactly how their information is being accessed, used and shared. The Arctera Insight Platform is able to capture data from over 130+ different content sources, and more than 280 AI policies help firms streamline compliance and adapt to evolving regulations.