Technical Information Sheet
Gaia RCM Expert (IIA FRANCE)

Gaia RCM Expert is an agent based on large language models (LLMs - Generative AI) intended for IIA FRANCE members and partners and accessible through the Gaia platform.

The main purpose of Gaia RCM Expert is to assist users in developing risk matrices, control frameworks and audit programs. To this end, it relies on a risk, control and audit test repository developed by IIA FRANCE.

In accordance with Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonized rules on artificial intelligence (the "AI Act"), Gaia RCM Expert falls within the category of limited-risk AI systems. Gaia RCM Expert does not fall within the high-risk system categories listed in Annex III and does not carry out any activities prohibited under Article 5.

Under Article 50 of the Regulation, these systems must comply with transparency obligations toward users when they interact directly with AI, in particular by clearly informing them that they are interacting with an AI system.

As deployer, IIA FRANCE keeps this documentation up to date and makes it available to the authorities upon request[1].

Making this documentation available to users also helps ensure that the system is accompanied by clear and appropriate information.

This technical information sheet covers Gaia RCM Expert only.

Separate technical information sheets cover these agents:

- *Gaia*
- *Gaia Lex*
- *Gaia Observation*
- *Gaia RCM*
- *Gaia Writer*

[1] In France, the authority responsible for supervising the transparency obligations laid down in Article 50 is the Directorate General for Competition Policy, Consumer Affairs and Fraud Control (DGCCRF).

Contents

1 Definitions

Artificial intelligence system (AI system): within the meaning of Article 3, point 1 of the AI Act, an AI system is "a system based on AI techniques capable, for a given set of human-defined objectives, of generating outputs such as predictions, recommendations or decisions influencing the environments with which they interact".

Deployer: under Article 3, point 4, any natural or legal person using an AI system under its authority, except where the system is used in the course of a purely personal, non-professional activity.

Provider: the entity that develops or places the AI model or system on the market.

General-purpose AI model (GPAI): within the meaning of Article 3, point 63, an AI model that can be used in a variety of applications for general purposes (e.g. GPT-4o). The obligations specific to GPAI models are set out in Article 53 of the Regulation and apply to the provider.

RAG (Retrieval-Augmented Generation): an approach that combines a language model with a structured document repository. Before generating a response, the system searches an internal corpus for the most relevant documents and then provides these materials to the model so that it can produce a contextualized response. This notably improves the reliability and traceability of responses (access to sources).

2 General identification of the Gaia RCM Expert AI system

System name	Gaia RCM Expert
AI model provider	**Azure OpenAI:** OpenAI models (gpt-5.4, gpt-5-mini, embeddings-3-large). Inference performed on *Microsoft Azure infrastructure (Europe)*
Deployer	IIA FRANCE
Designated person responsible	Jean Loup Grosse – Head of Systems and Organization, IIA FRANCE
Hosting	Gaia infrastructure (servers, data): **OVH - France**
Go-live date	December 2025
Current version	Gaia 2.4.1 (May 2026)

3 Description of the Gaia RCM Expert agent

Gaia RCM Expert is an AI agent accessible through IIA FRANCE's Gaia platform.
It assists the user in preparing the work program and the risk/control/audit test matrix (RCM), based on the user's audit context.

- An initial conversational phase guides the user in defining the audit context.
- Once the audit context and parameters have been defined, the user can proceed to the generation phase (using the "Generate the RCM" button).
- A reasoning model then analyzes the audit context and queries a previously established reference database of risks, controls and tests using various strategies. It may also propose additional items based on its own capabilities.
- At the end of this phase (5 to 10 minutes), the following are submitted for the user's review:
    - A list of audit objectives
    - For each audit objective, a list of risks, controls and tests:
        - Either taken from or inspired by the IIA FRANCE reference database and adapted to the audit context
        - Or generated by the language model without reference to the IIA FRANCE database
    - An executive summary
- The user may assess or modify the various elements of the matrix and export the result to Excel.

4 Impact, risks and mitigation

Users concerned: IIA FRANCE members and partners with access to the Gaia RCM Expert agent.

4.1 Expected positive impacts

The deployment of the Gaia RCM Expert agent is intended to generate several positive impacts, including:

- Development of user skills: through structured, sourced and contextualized responses, Gaia RCM Expert supports users in understanding and adopting professional good practices.
- Improved productivity: Gaia RCM Expert frees up time for analysis and human validation.
- Support for quality and compliance: generation, validation and source-access features reinforce methodological rigor and reduce omissions or inaccurate wording.

4.2 Risks/Potential negative impacts

Despite these benefits, several risks must be taken into account:

- Risk of factual or interpretive errors: some responses may be incomplete, approximate or taken out of context, particularly where wording is ambiguous or because of limitations inherent in AI models.
- Risk of reliance without human verification: users may be tempted to regard responses as accurate without validation, which could lead to erroneous decisions or the dissemination of incorrect information.
- Risk of poor wording or omission of essential elements: certain key components may be missing or poorly structured.
- Risk of overreliance: intensive use could reduce critical thinking or direct consultation of the reference materials.
- Residual risk of bias or hallucinations: as with any conversational model, the production of incorrect information remains possible, even with an effective RAG system.
- Risk of inappropriate use: the agents could be used for purposes inconsistent with their intended use (for example, outside the intended professional or regulatory framework), or to produce inappropriate or unauthorized content.

4.3 Mitigation measures

To limit these risks, several technical, organizational and methodological measures have been implemented:

- Use of RAG: part of the responses is based on a database of risks, controls and tests derived from a validated document corpus (IPPF, standards and IIA FRANCE documents), limiting hallucinations and improving reliability. A reformulation indicator *(Adaptation)* is displayed for each row of the matrix relative to the reference database.
- Disclaimers: a disclaimer appears at the start of the conversation and in the RCM generation interface, and is added to every Excel export.
- Instructions: the system prompts are designed to limit risks from the outset. They govern the behavior of Gaia RCM Expert by specializing it in audit and internal control topics, defining the expected types of responses, and incorporating professional good practices directly into its core instructions.
- Guided output: generation of the audit context follows a structured process before the RCM is generated.
- User training and awareness: prompt-writing assistance is integrated directly into Gaia. IIA FRANCE also offers Gaia training sessions.

- Human review and oversight: the user is invited to review, assess and modify the various elements of the matrix before exporting it to Excel. The reporting features *(Like/Dislike) allow* problematic responses to be flagged for analysis and continuous improvement.

5 Confidentiality and data protection

5.1 No model training

Inference for the models used is performed in Europe on Azure infrastructure (Azure OpenAI). Prompts (questions) and completions (responses) are not used to improve or train the OpenAI or Microsoft models[2].

5.2 Sovereign hosting

The Gaia RCM Expert infrastructure (servers and document repository) is hosted in France in an OVH environment dedicated to IIA FRANCE.

5.3 Data processing and storage

Gaia RCM Expert does not store questions or its complete responses[3].
The complete matrix and the executive summary are not stored. The conversation (generation of the audit context) and the RCM are stored locally in the user's web browser.

The user can clear the conversation from the browser's local storage at any time (using the *"New conversation"* and *"Delete all conversations saved locally" buttons).*

The following are retained anonymously for the purpose of improving the platform:

- a brief intent, expressed in a few words and inferred from the audit objective (for example, "procurement and inventory audit");
- the relevant macro-process and industry (for example, "Deliver Physical Products" and "Manufacturing");
- a condensed version, stripped of any sensitive information, of risks, controls and tests not sourced from the IIA FRANCE database (cells in italics in the RCM);
- the risk reformulation indicator relative to the reference database;
- any feedback and ratings.

[2] https://learn.microsoft.com/en-us/azure/ai-foundry/responsible-ai/openai/data-privacy
[3] With the potential exception of inappropriate content; see "Detection and prevention of inappropriate content".

5.4 Personal and confidential data

No personal data (name, email address, etc.) is stored in Gaia RCM Expert.
Only a pseudonymized identifier, generated from the user's login information using a hashing function combined with encryption, is associated with each request. It is used solely for aggregated usage statistics (typically the number of unique platform users).

In accordance with Article 4 of the GDPR, this identifier constitutes pseudonymized personal data. As it stands, it does not allow third parties to identify an individual directly without additional information.
In accordance with the preceding section, "Data processing and storage," any personal or confidential data submitted by the user during use (through prompts) is not stored in Gaia[4].

5.5 Detection and prevention of inappropriate content

A system for analyzing requests (prompts) and responses uses classification models to detect inappropriate content and prevent its generation. These models cover four categories defined by Azure: hate, sexual content, violence and self-harm.
If the system is triggered, response generation is blocked and Gaia RCM Expert displays an error. In the event of repeated detections of abuse, IIA FRANCE may suspend or prohibit access to Gaia RCM Expert for the identifier concerned.

In addition, where the classification model detects inappropriate content, Microsoft may conduct an additional review:

- First, automatically using a more advanced model (LLM)
- Then, where applicable, by a human operator if the model determines that the automated review is insufficient.

The automated review complies with the principles set out in Sections 5.1 (no model training), 5.3 (no storage of prompt/response data) and 5.4 (personal data).
If the case is escalated for human review, the prompt and response are stored by Microsoft for a maximum of 30 days in the Europe region. Microsoft personnel authorized to analyze inappropriate content are located in the European Economic Area.
Except for potentially inappropriate content, prompts and responses are not stored by OpenAI/Microsoft (inference via the *Response API without State Storage)*[5]

[4] With the potential exception of inappropriate content; see "Detection and prevention of inappropriate content".
[5] https://learn.microsoft.com/en-us/azure/ai-foundry/responsible-ai/openai/data-privacy

5.6 Response evaluation

Gaia RCM Expert includes a feature allowing users to evaluate the quality of RCM elements through *Like/Dislike/Needs review* buttons placed beneath each cell of the matrix.
In addition, when exporting to Excel, the user is invited to rate their experience by providing a score from 1 to 5 and an optional comment.
This information is stored on the platform when the Excel export is performed, for statistical analysis and to support future developments of Gaia RCM Expert. It is not stored if the RCM is not exported to Excel.

The conversational phase for guided generation of the audit context includes a *Like/Dislike* reporting feature.
When a user clicks one of these buttons, they indicate the reason and may leave a free-text comment. This comment is stored in the Gaia infrastructure to enable human review by IIA FRANCE. The corresponding conversation is neither transmitted nor stored.

5.7 Usage metrics and logs

The following anonymized metrics are collected and, where applicable, regularly analyzed by the designated IIA FRANCE representative:

- Number of daily uses of Gaia RCM Expert
- Number of unique users
- Number of accesses to source documents
- Number of accesses to IIA FRANCE training courses
- Number of response reports
- Number and types of errors encountered
- Inappropriate-content detections
- Average rating
- Anonymized comments
- Performance of the RAG system and reasoning LLM

To enable these analyses, the following events are logged and may be associated with the non-reversible pseudonymized identifier:

- Use of the Gaia RCM Expert agent, associated intent (summary of no more than five words, for example, “procurement and inventory audit”), and associated process and industry (closed lists, for example, “Deliver Physical Products” and “Manufacturing”)
- Viewing a source document
- Redirection to an IIA FRANCE training course
- Use of the *Like / Dislike / Needs review* features
- Error while generating a response

- Condensed version, stripped of any sensitive information, of risks, controls and tests not sourced from the IIA FRANCE database (cells in italics)
- List of risk, control and test items from the IIA FRANCE database used to build the response.
- Reformulation indicator for each risk relative to the IIA FRANCE database.

Prompts, responses, the audit context and the RCM are never stored in these usage logs.

6 Obligations relating to GPAI models (Chapter V of Regulation (EU) 2024/1689)

The language and embedding models used by Gaia RCM Expert (gpt-5-4, gpt-5-mini and embeddings-3-large) are general-purpose AI models (GPAI models) within the meaning of Article 3, point 63 of Regulation (EU) 2024/1689.

IIA FRANCE is not a GPAI model provider within the meaning of the Regulation. It acts solely as a deployer, using models provided by OpenAI and distributed through Microsoft Azure.

In accordance with Articles 52 to 55, the preparation and publication of the "model summary" (GPAI information sheet) are the provider's responsibility. These information sheets will be appended to this document as soon as they are officially published.

In the meantime, the public documentation available for each model is listed below:

Model	Technical documentation	System Card
gpt-5.4	https://platform.openai.com/docs/models/gpt-5.4	https://openai.com/index/gpt-5-4-thinking-system-card/
gpt-5-mini	https://learn.microsoft.com/en-us/azure/foundry/foundry-models/concepts/models-sold-directly-by-azure?tabs=global-standard&pivots=azure-openai#gpt-5	https://openai.com/index/gpt-5-system-card/
embeddings-3-large	https://openai.com/index/new-embedding-models-and-api-updates	Not applicable (embedding model)