




Kaspersky TIARA

LLMs can act as intelligent agents, interacting with external systems by generating requests to their APIs. These agents can perform complex tasks, such as searching for and purchasing products from an online store based on user preferences.

The risk of hallucinations is particularly high when an LLM's input data pertains to entities and facts that are poorly represented in the model's training dataset, or when the input contains information that misleads the LLM, either accidentally or intentionally.



The amount of available domain-specific knowledge often exceeds the capacity of an LLM's context window. However, answering a user's specific query typically requires only a small fraction of this knowledge. TIARA's agentic graph RAG enables LLMs to generate responses without context overflow or unnecessary computational overhead.

Technology for integrating industry-specific knowledge into large language models

Let your expertise talk like a human.

What's the goal of Kaspersky TIARA?

Kaspersky Trusted Intelligent Agents for Retrieval and Analysis (TIARA) is a comprehensive technology that enhances the accuracy and reliability of responses from large language models (LLMs) where it matters most — in real-world industry applications.

Knowledge vs. hallucinations

Due to their probabilistic nature, large language models can generate hallucinations, i.e., false or nonsensical information presented as facts.

Kaspersky TIARA builds upon Retrieval-Augmented Generation (RAG) technology to equip LLMs with real, domain-specific knowledge. This minimizes hallucinations and erroneous actions when operating as AI agents, as well as potential harm to businesses.

Graph as a knowledge representation

Kaspersky TIARA transforms large collections of domain-specific documents, such as encyclopedias, scientific publications, reference manuals, and regulatory codes, into a unified knowledge graph. The graph comprises logical document elements (chapters, sections, paragraphs, figures, tables) as well as extracted entities and facts.

TIARA supports static documents in various formats (PDF, DOCX, HTML) along with image, video, and audio files.

Agentic graph RAG

Kaspersky TIARA is built on AI agents that it leverages to search the knowledge graph for information relevant to the user's query. The agents automatically develop optimal retrieval strategies that take into account the graph's structure, as well as the vocabulary and semantics of the texts.


Cyberattacks on AI models are carried out by inserting specially crafted “corrupted” data fragments into input queries or knowledge source documents. These attacks cannot be detected by conventional antivirus software.

Accuracy and security control

Kaspersky TIARA incorporates mechanisms for monitoring response generation. These mechanisms analyze the user's query, the retrieved knowledge, and each step of the LLM generation process. This ensures the accuracy and completeness of responses, and provides protection against AI-targeted cyberattacks, such as poisoning and jailbreaks.

Architecture of Kaspersky TIARA

Examples of domains that may benefit from Kaspersky TIARA



Is Kaspersky TIARA suitable for your organization?

1 On-premises deployment

TIARA does not rely on cloud services, operating entirely within the customer's infrastructure and thereby ensuring the confidentiality of processed data.

2 Scalability

TIARA is built on high-performance AI models and data indexing tools, enabling it to process gigabytes of documents on consumer-grade hardware.

3 Learning curve

TIARA does not require AI specialists for maintenance. The selection of AI algorithms and their parameters is performed automatically and takes into account the specifics of the domain and user feedback.



Learn more at
ai-cert.kaspersky.ru

kaspersky

www.kaspersky.com

© 2026 AO Kaspersky Lab.
Registered trademarks and service marks are the property of their respective owners.