Corelight AI Trust FAQs

Welcome to Corelight's AI Trust page. Here Corelight addresses questions regarding the use of Corelight's AI Features. As used in this document "**AI Features**" refers to artificial intelligence (AI) technology that uses inputs to predict outcomes or that is capable of generating new content included as part of Corelight's offerings.

This document does not amend or form part of your contract terms with Corelight and the information contained herein may evolve over time. For more information about Corelight's security, privacy and compliance programs, please visit the Corelight Trust Center (available at www.corelight.com/trust-center).

Corelight Investigator

Corelight Investigator ("**Investigator**") leverages AI to enhance the capabilities of its platform and provide users with the tools to streamline network security investigations.

"**AI Assistance**" is a suite of Investigator AI Features that use large language models (LLMs) to understand network data and provide intelligent assistance. This suite of Investigator AI Features comprises:

- Rule Summary & Impact
- Alert Insights & Payload Summary

How do Investigator's AI Assistance capabilities work?

- **Rule Summary & Impact**

  Investigator's Rule Summary & Impact is an LLM-powered assistant that simplifies complex threat detection logic and seamlessly integrates into a customer's existing Investigator detection workflow. To start, Investigator's Rule Summary & Impact helps improve users understanding of why an alert was generated by automatically displaying a plain English "Rule Description" of the underlying Suricata rule. Users can then interact with Investigator's pre-populated prompts to receive additional explanations regarding alert meaning and next steps to investigate the alert.

- **Alert Insights & Payload Summary**

  Investigator's Alert Insights & Payload Summary is an LLM-powered assistant that contextualizes an alert's supporting logs and payload to accelerate a customer's alert investigation. Using Investigator's alert sidebar panel, users can click "Analyze Activities" and Investigator will generate "Alert Connection Insights" based on the logs associated with the alert. The sidebar also automatically displays a plain English "Payload Summary" along with extracting key findings from the packet's payload data.

How does a user know when AI Features are part of the Investigator experience?

Corelight is working to ensure that there are consistent visual indicators across a user's Investigator experience, similar to this icon, to specify where AI Features are being used. Users may also see other icons, product notifications, and notices indicating the product or feature is AI-assisted.

Which LLMs are being used for Investigator's AI Assistance?

Corelight employs best-in-class third-party hosted LLMs. As of the document's publication date, Investigator's AI Assistance uses OpenAI's GPT series of models, accessed via API.

Please refer to Corelight's list of data subprocessors for more information on our third-party hosted LLM providers.

How does Investigator use data when users engage with AI Assistance capabilities?

- **Rule Summary & Impact**: When users engage with Rule Summary & Impact, ***no*** customer data is used. Rule Summary & Impact processes Corelight-authored prompts along with Corelight-provided Suricata rules to provide the outputs.

- **Alert Insights & Payload Summary**: When users click on "Analyze Activities", customer data (e.g., session logs) are submitted to and processed by the third-party LLM to generate "Alert Connection Insights". Packet payload data is submitted to and processed by the third-party LLM to provide the "Payload Summary".

Does use of AI Assistance result in the transfer of any data outside of Investigator?

When using Investigator's AI Assistance capabilities, certain data (as further detailed below) is transferred outside of Investigator to a third-party LLM provider (e.g., OpenAI) in order to generate a response. Each data request is sent to the third-party LLM provider individually, over an SSL-encrypted service, to process, and send back to Investigator.

- **Rule Summary & Impact**: transfers are limited to publicly available data (i.e., Corelight-authored prompts along with Corelight-provided prompts).
- **Alert Insights & Payload Summary**: transfers involve select categories of customer data (i.e., session logs and payload data).

When Corelight partners with third-party LLM providers such as OpenAI, how is data protected?

When working with third-party LLM providers, such as OpenAI, Corelight and its LLM providers agree on appropriate data protections to safeguard the data privacy and security of Corelight's customers and users. The LLM provider Corelight uses does not retain inputs or outputs or use them to improve their services.

Does the third-party LLM provider store data?

No, Corelight's LLM provider does not store the data a user submits, or the responses users receive. Data is immediately deleted by the LLM provider after processing (zero data retention).

Does the third-party LLM provider use data sent from Investigator for AI model training and improvement?

No, the data submitted and responses received from AI Assistance are not used to train, fine-tune or improve any AI model(s) or services.

Does Investigator's AI Assistance use customer data to serve other customers?

No, any customer data submitted and the responses received are used solely for the benefit of that specific customer's experience. Customer data is not used to train models across customers or shared between customers.

What is the architecture underlying Investigator's AI Assistance capabilities?

Investigator leverages secure data ingestion and processing, supported by backend services for metadata enrichment and indexing. Investigator's AI Assistance leverages OpenAI's GPT series of models, accessed via API. Prompts are written internally by Corelight and pre-populated within Investigator's web-based UI. Likewise, users interact with AI-generated outputs via Investigator's UI. Users do not provide inputs or interact directly with the LLM.

Information on how OpenAI trains its models is available here. Given that third-party foundational LLMs are trained on publicly available data, Corelight benefits from its open-core model contributing to such training.