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 Generative AI technology included as part of Corelight's offerings.
- "**Generative AI**" refers to machine learning or artificial intelligence systems, including large language models, that generate content in response to inputs (but excluding systems used solely for analytical purposes such as classification, prediction or recommendation).

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).

How does Corelight use Generative AI in its offerings?

Corelight's generally available AI Features are included as part of Corelight's Investigator SaaS offering ("Investigator"). Currently, Investigator's AI Features include:

- "**AI Assistance**" uses large language models (LLMs) to understand network data and provide intelligent assistance. This suite of Investigator AI Features comprises:
    - Rule Summary & Context
    - Alert Insights & Payload Summary
- "**Agentic Triage**" is a suite of Investigator AI Features that employ predefined workflows to perform initial triage, validation and correlation of Corelight detections.

How do Investigator's AI Assistance capabilities work?

- **Rule Summary & Context**

  Investigator's Rule Summary & Context 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 & Context 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 do Investigator's Agentic Triage capabilities work?

Investigator's Agentic Triage agents work automatically in the background, identifying high risk network entities, performing triage and correlative analysis and summarizing findings and evidence for review and action by human analysts. Agentic Triage outputs are incorporated directly into Investigator's analysis workflows. Notably, the triage confidence score is located in the entity panel in the Dashboard view, where triaged entities are categorized as "Likely Malicious", "Suspicious" or "Likely Benign". When the user opens an entity detail view, Agentic Triage outputs are again incorporated into the workflow and clearly marked with

the  icon. The entity detail view includes triage agent analysis summary, detection and entity summary, key findings and recommended actions.

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

Corelight is working to ensure that there are consistent indicators across a user's Investigator experience, currently similar to this  icon, to specify where AI Features are being used. Corelight will identify AI Features within the Corelight offerings themselves or in the documentation.

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

Corelight employs a best-in-class third-party hosted LLM. As of this document's publication date, Investigator's AI Features solely use 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 provider.

How does Investigator use data when users engage with AI Features?

- **Rule Summary & Context**: When users engage with Rule Summary & Context, ***no*** customer data is used. Rule Summary & Context 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".
- **Agentic Triage**: Agentic Triage runs on a schedule with no additional human intervention. When the workflows are executed, the specific customer data needed to investigate new detections is retrieved, then submitted to, and processed by the third-party LLM to generate analysis summaries, confidence scores, key findings and recommended actions. When users click on an "AI Insights" or interact with any Agentic Triage output, no additional customer data is used.

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

When using Investigator's AI Features, 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 TLS-encrypted service, to process, and send back to Investigator.

- **Rule Summary & Context**: transfers are limited to publicly available data (i.e., Corelight-authored prompts along with Corelight-provided Suricata rules).
- **Alert Insights & Payload Summary**: transfers involve select categories of customer data (i.e., session logs and payload data).
- **Agentic Triage**: transfers include select categories of customer data, including aggregated summaries of threat detections, network connections and protocol transactions.

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

When working with its 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 (i.e., OpenAI) that Corelight uses does not retain inputs or outputs or use them to improve their services.

Does the third-party LLM provider (i.e., OpenAI) store data?

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

Does Corelight use customer data to train models used by Investigator's AI Features?

Investigator currently leverages pre-trained LLMs provided by OpenAI. No customer data is used to train or fine-tune these models.

Does Corelight use customer queries or prompts to train models used by Investigator's AI Features?

Investigator currently leverages pre-trained LLMs provided by OpenAI. Customer queries and prompts are not used to train or fine-tune these models.

Do Investigator's AI Features 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 Features?

Investigator leverages secure data ingestion and processing, supported by backend services for metadata enrichment and indexing. Investigator's AI Features leverage 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.

Can customers disable or opt out of Investigator's AI Features?

Yes. Corelight is committed to providing customers with choice around the use of AI Features.

For new customers purchasing Investigator on or after April 1, 2026, Investigator's AI Features will be enabled by default. Investigator admins can manage (disable or enable) AI Features via Investigator's integration settings menu. Customers can learn more about managing AI Features within Investigator in Corelight's documentation.

* *For customers that purchased Investigator prior to April 1, 2026, AI Features are disabled by default. Please contact Corelight Support for configuration assistance*.