opentext™

GUIDE

A CISO's guide to an AI-enhanced SOC

How AI models, enhanced pipelines, and the MITRE ATT&CK framework can help CISOs scale detection, reduce noise, and improve SOC effectiveness without increasing headcount

Contents

CISOs face an operational paradox: more threats, more alerts, and more responsibilities—but no more staff. At the same time, executive expectations (and industry regulations) are rising. Security teams must detect faster, respond smarter, and prove they can manage risk proactively.

Enter AI.

Large language models (LLMs) and enhanced retrieval-augmented generation (RAG) workflows are reshaping how security operations centers (SOCs) investigate alerts. By linking high-volume telemetry to the MITRE ATT&CK framework with explainable, auditable reasoning, these AI-enhanced pipelines help security teams gain context, clarity, and confidence without adding more data collection solutions or people. This guide is for CISOs who want to make AI practical, trustworthy, and actionable in their SOC strategy. It introduces the key concepts behind AI-enhanced threat hunting, how OpenText™ Core Threat Detection & Response implements them, and why it matters for organizations striving to do more with what they already have.

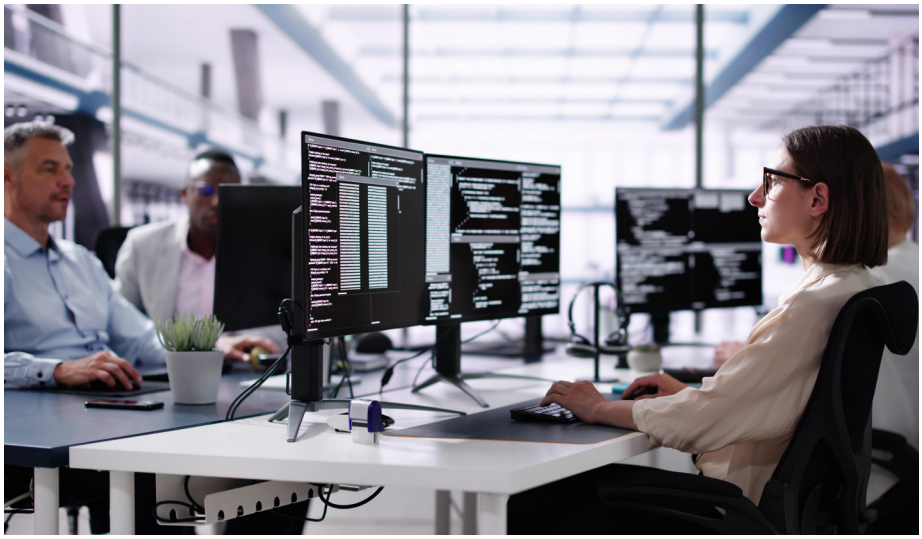

Addressing the elephant in the room: Can you trust AI in the SOC?

There's no shortage of uncertainty and doubt around AI in cybersecurity. It's often portrayed as a black box, a job killer, or a liability waiting to happen. But correct implementation can drastically reduce these concerns. **AI won't replace your SOC**: It removes repetitive noise, not critical thinking. Analysts remain in control.

- **LLMs do hallucinate**: but not when grounded: Enhanced RAG ensures LLMs receive structured, relevant input and produce traceable, auditable results.
- **You need explainability, not opacity**: Every output includes human-readable reasoning supporting compliance, reporting, and board-level visibility.
- **Automation is always optional**: AI augments your workflows. It doesn't execute action unless you define it.

At OpenText, we've engineered our system to address these concerns. OpenText Core Threat Detection & Response combines domain-specific modeling, scoring, and retrieval strategies to make AI trustworthy, scalable, and useful. That's not theory it's the product.

The CISO mandate: More threats, same resources

Cybersecurity leaders are expected to deliver better outcomes; faster detection, earlier intervention, and fewer missed threats, while navigating a constrained operating environment. That includes overworked SOC teams, an ever-expanding threat surface, and escalating executive scrutiny.

Key pressure points include:

- **SOC fatigue and burnout**: Alert triage is unsustainable. Teams are drowning in signals, many of which lack context or relevance.
- **Detection complexity**: Point solutions generate fragmented insights. Threats often span tools, formats, and timelines.
- **Framework expectations**: Boards and regulators increasingly expect detection mapped to MITRE ATT&CK or similar standards. Manual mapping is impractical at scale.

In this environment, more tooling isn't the answer. More insight is. CISOs must adopt solutions that turn raw telemetry already being collected into structured, actionable intelligence without adding headcount or introducing risk. AI, properly implemented, offers that leverage.

Moving beyond the hype: Real use of LLMs

CISOs are bombarded with AI claims but most lacked operational depth. Early efforts fell short:

- **Manual tagging** couldn't scale.
- **ML classifiers** lacked high-quality, labeled data.
- **Generic LLMs** hallucinated or misinterpreted signal data.
- The turning point: Enhanced retrieval-augmented generation (RAG).
- **RAG** connects LLMs to structured, curated knowledge, in this case, MITRE ATT&CK.
- **Enrichment** gives alerts more context (e.g., threat intel on process names).

What's different here is our method.

Prompt engineering and chunking is where our data science brilliance comes in. When combined, the model receives the right information, in the right format, yielding superior results, drastically reducing hallucinations. This fusion of structure and language makes LLMs useful, not just novel. They shift from content generators to context engines.

For a deeper look at how these models are tuned, trained, and evaluated within OpenText, download the white paper, Inside the Data Science of OpenText Core Threat Detection & Response. It includes details on probabilistic modeling, anomaly detection baselines, peer clustering, dimensionality reduction, and LLM validation techniques.

How OpenText Core Threat Detection & Response empowers teams with AI

OpenText has invested years refining its threat detection stack. With OpenText Core Threat Detection & Response, the AI is built into the pipeline, not bolted on.

Capabilities include:

- **Input enrichment**: Adds threat intelligence and historical behavior to raw alerts.
- **Semantic chunking**: Groups MITRE ATT&CK content meaningfully to improve retrieval fidelity.
- **Scoring engine**:
  - *Frequency:* How often a technique shows up in signal clusters.
  - *Stage:* Where it fits in the attack lifecycle.
  - *External validation:* Whether it appears in threat reports.
- **LLM-driven reasoning**: Every TTP match is accompanied by a clear explanation turning SOC outcomes into evidence.
- **Evaluation with LLM-as-judge**: Final mappings are verified by larger LLMs to ensure consistency and validity.

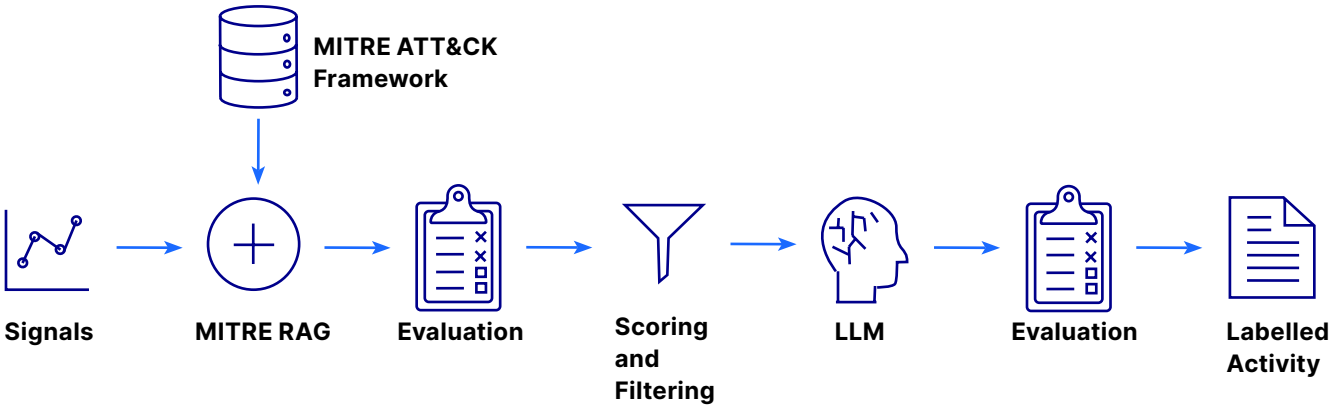

Together, these components deliver rich, MITRE ATT&CK-aligned insights without human bottlenecks.

A real-world scenario:

A service account begins behaving abnormally logging in via remote sessions, launching atypical processes, and accessing unusual resources.

Old approach using traditional detection methods

- A flood of individual alerts
- Analysts manually review behavior
- Time-consuming correlation

New approach with OpenText Core Threat Detection & Response

- Signals aggregated and enriched by OpenText Core Threat Detection & Response
- System retrieves likely MITRE ATT&CK techniques (e.g., T1110 Brute Force, T1021 Command and Control)
- Scoring ranks the threat priority
- LLM delivers a summary: “The service account is likely compromised and being used to laterally move through the environment via remote services.”

The analyst doesn’t start with raw data. They start from a narrative and context, accelerating time to decision.

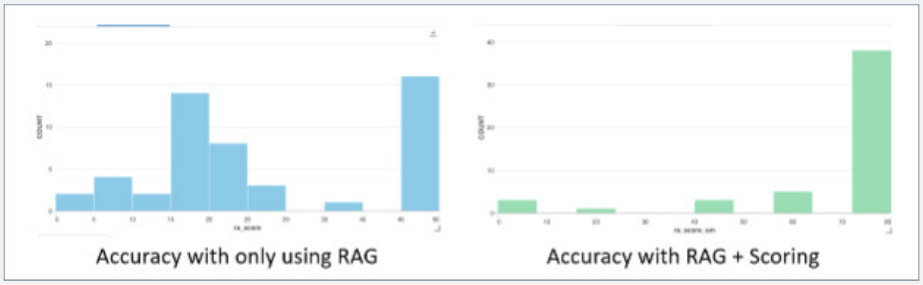

Resources

See a demo ›

Read the technical white paper ›

View the product page ›

Take a self-guided tour ›

Contact us ›

A scalable path forward

You can't hire your way out of the detection gap. But you can build a smarter, more scalable pipeline.

OpenText Core Threat Detection & Response enables CISOs to:

- Map more threats to MITRE ATT&CK with less human effort.
- Reduce alert fatigue and improve investigation outcomes.
- Empower junior analysts with senior-level context.
- Create an auditable, trustworthy AI detection layer.

This is the future of cyber defense: human-led, AI-accelerated, MITRE ATT&CK-aligned. It's time to get started.

Want to see the technical depth behind this guide? Read the companion white paper, Inside the Data Science of OpenText Core Threat Detection & Response, for a full breakdown of our modeling, scoring, and LLM integration methods.

opentext™