AI-Powered SOC: Cisco Sec-8b and AIPOD in Action

Daria Margherita Maggi
Consulting Engineer
Cisco

Stefano Gioia
EMEA Solution Architect
Cisco

BRKAI-1244

Webex App

Questions?

Use Webex App to chat with the speaker after the session

How

1. Find this session in the Cisco Events App
2. Click “Join the Discussion”
3. Install the Webex App or go directly to the Webex space
4. Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until February 27, 2026.

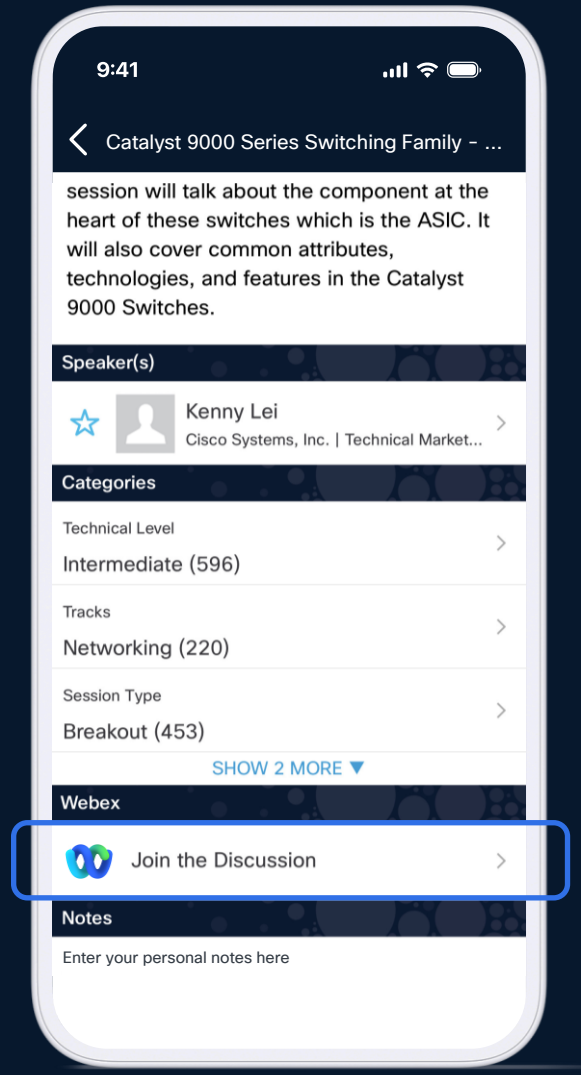

Daria Margherita Maggi

Adversarial Machine Learning with generative AI researcher (before it was cool :))

CX Security Consulting Engineer, specialized in military and highly regulated sectors.

Cat mom.

cisco

Stefano Gioia

25+ experience in Industry:
Ericsson, Oracle, Cisco since 2017

Observability, AI and
Virtualization expert.

BBQ Judge

Agenda

01 Introduction

02 LLMs and SOC troubles

03 Meet the Sec-8B Model

04 Sec-8B Model Use Cases

05 Meet the Cisco AI POD

06 Sec-8B in Action: Demo

07 Closing and Q&A

”Safe Harbor” Statement
This session is not...

Level 200, 300

It is Level 100 so it will be introducing concepts and practices.

Big Model Discussion

Mostly on-premise or air-gapped scenarios, with limited computing power – small language models only.

Vendor comparison

We are not comparing us versus other.

What you will get from this session …

How has AI transformed the way SOC manage and respond to threats.

How you can accelerate SOC process with Cisco Foundation AI model(s).

Cisco AI POD as the perfect platform to host Cisco Sec-8B Security Model.

LLMs and SOC troubles

AI has changed the threat scope

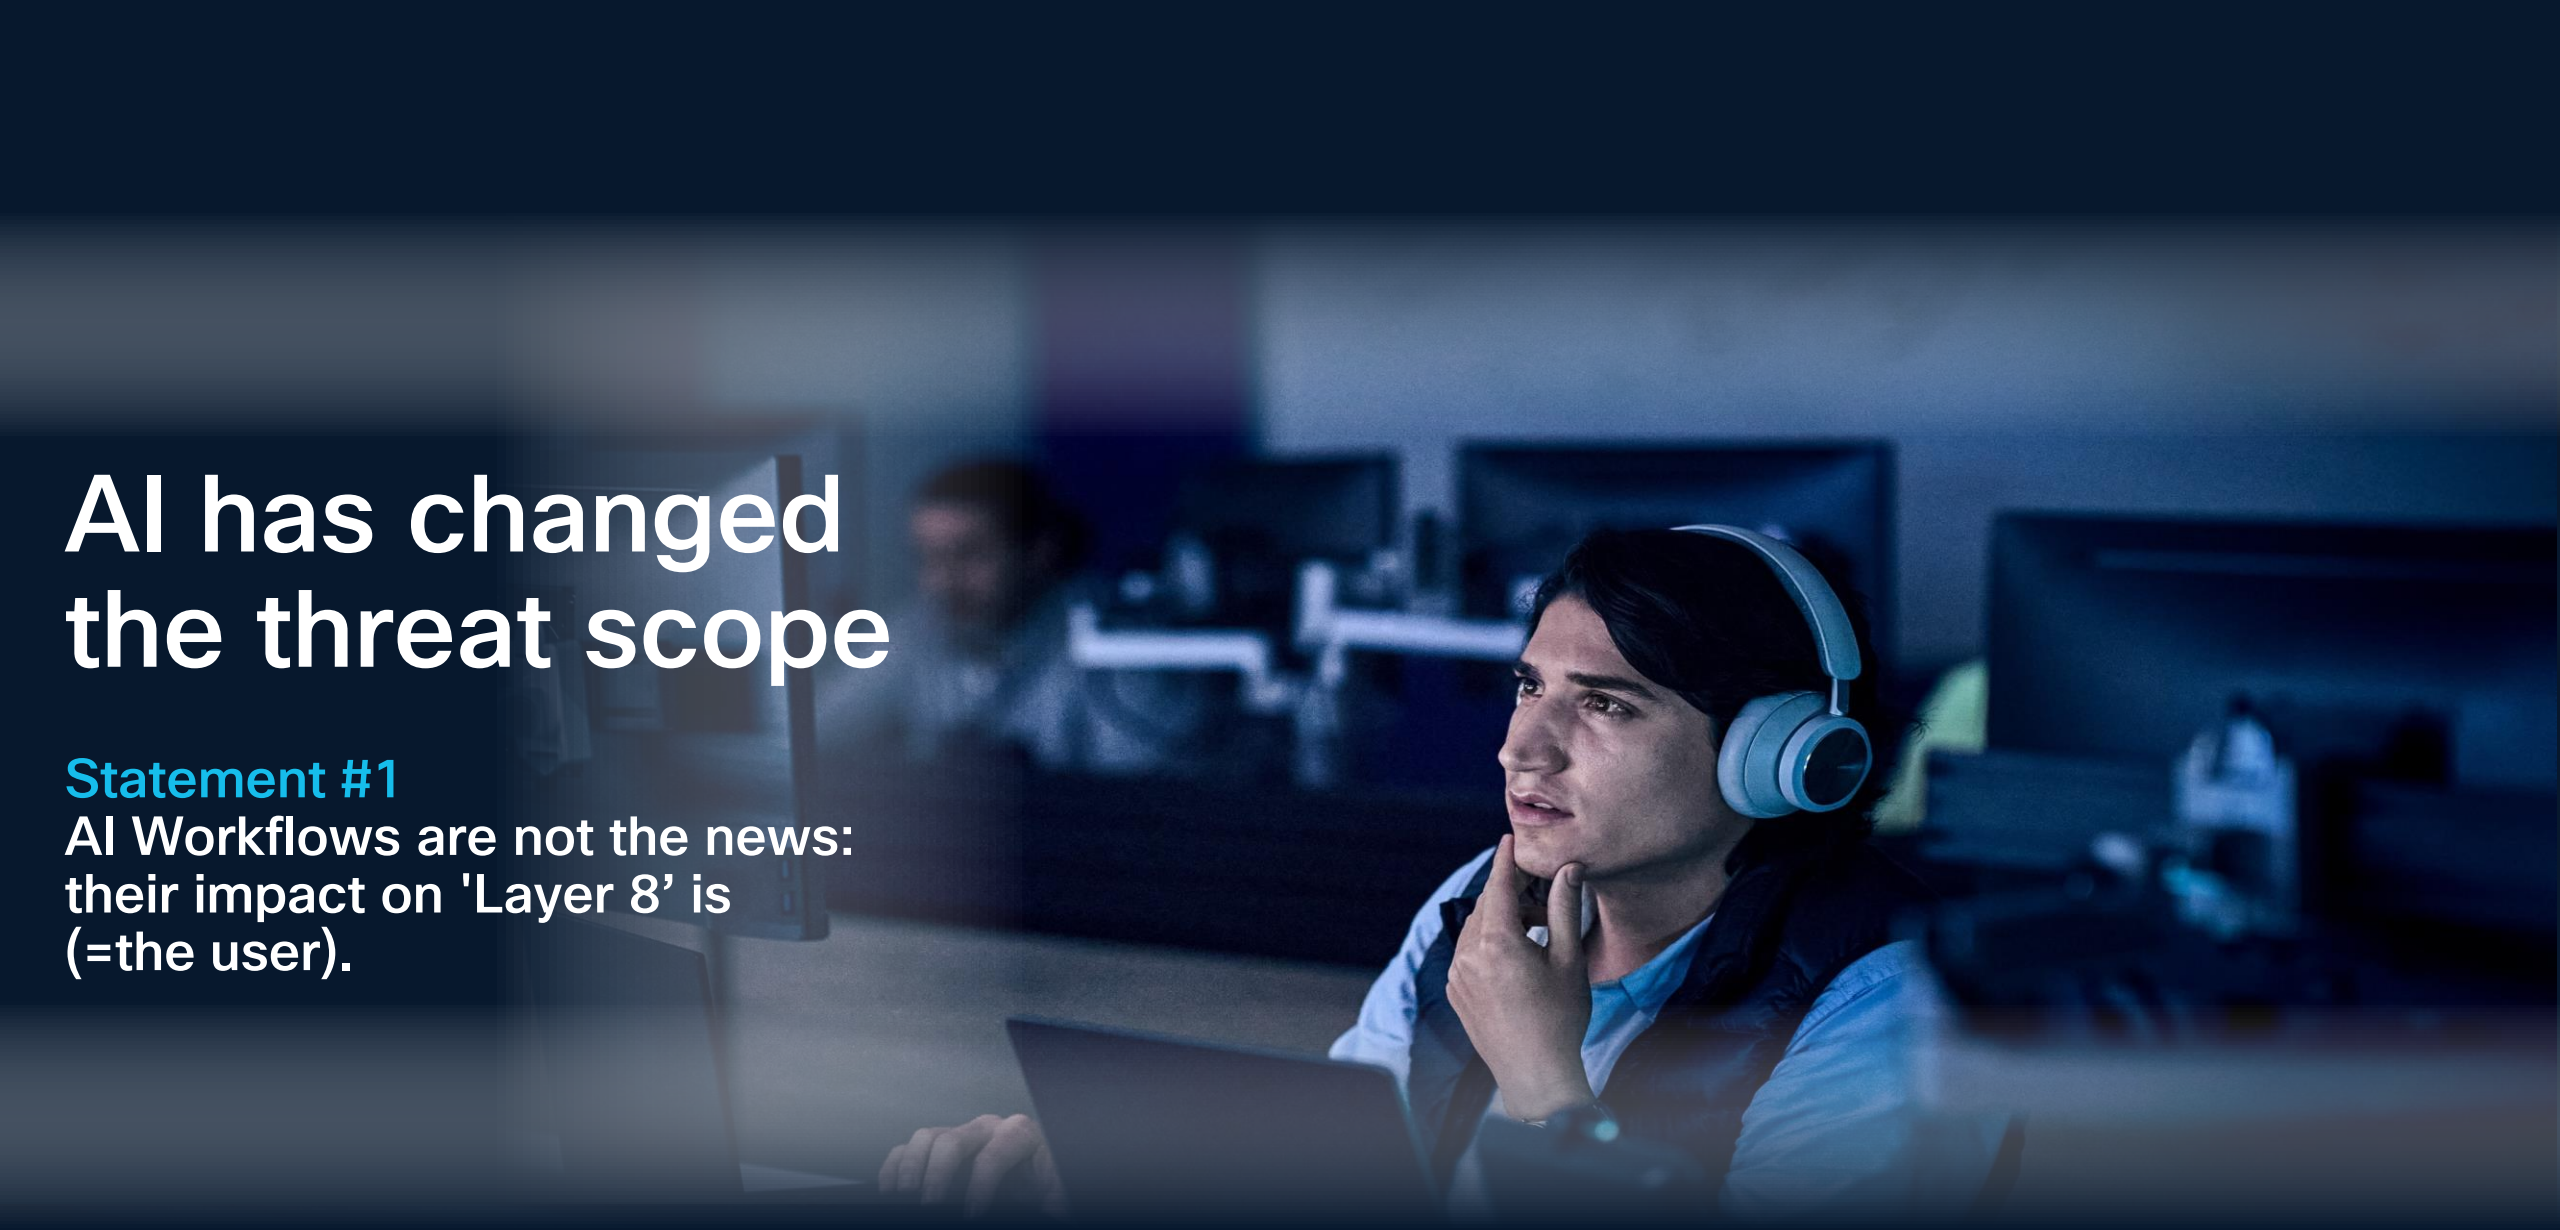

Statement #1

AI Workflows are not the news: their impact on 'Layer 8' is (=the user).

cisco

AI has changed the threat scope

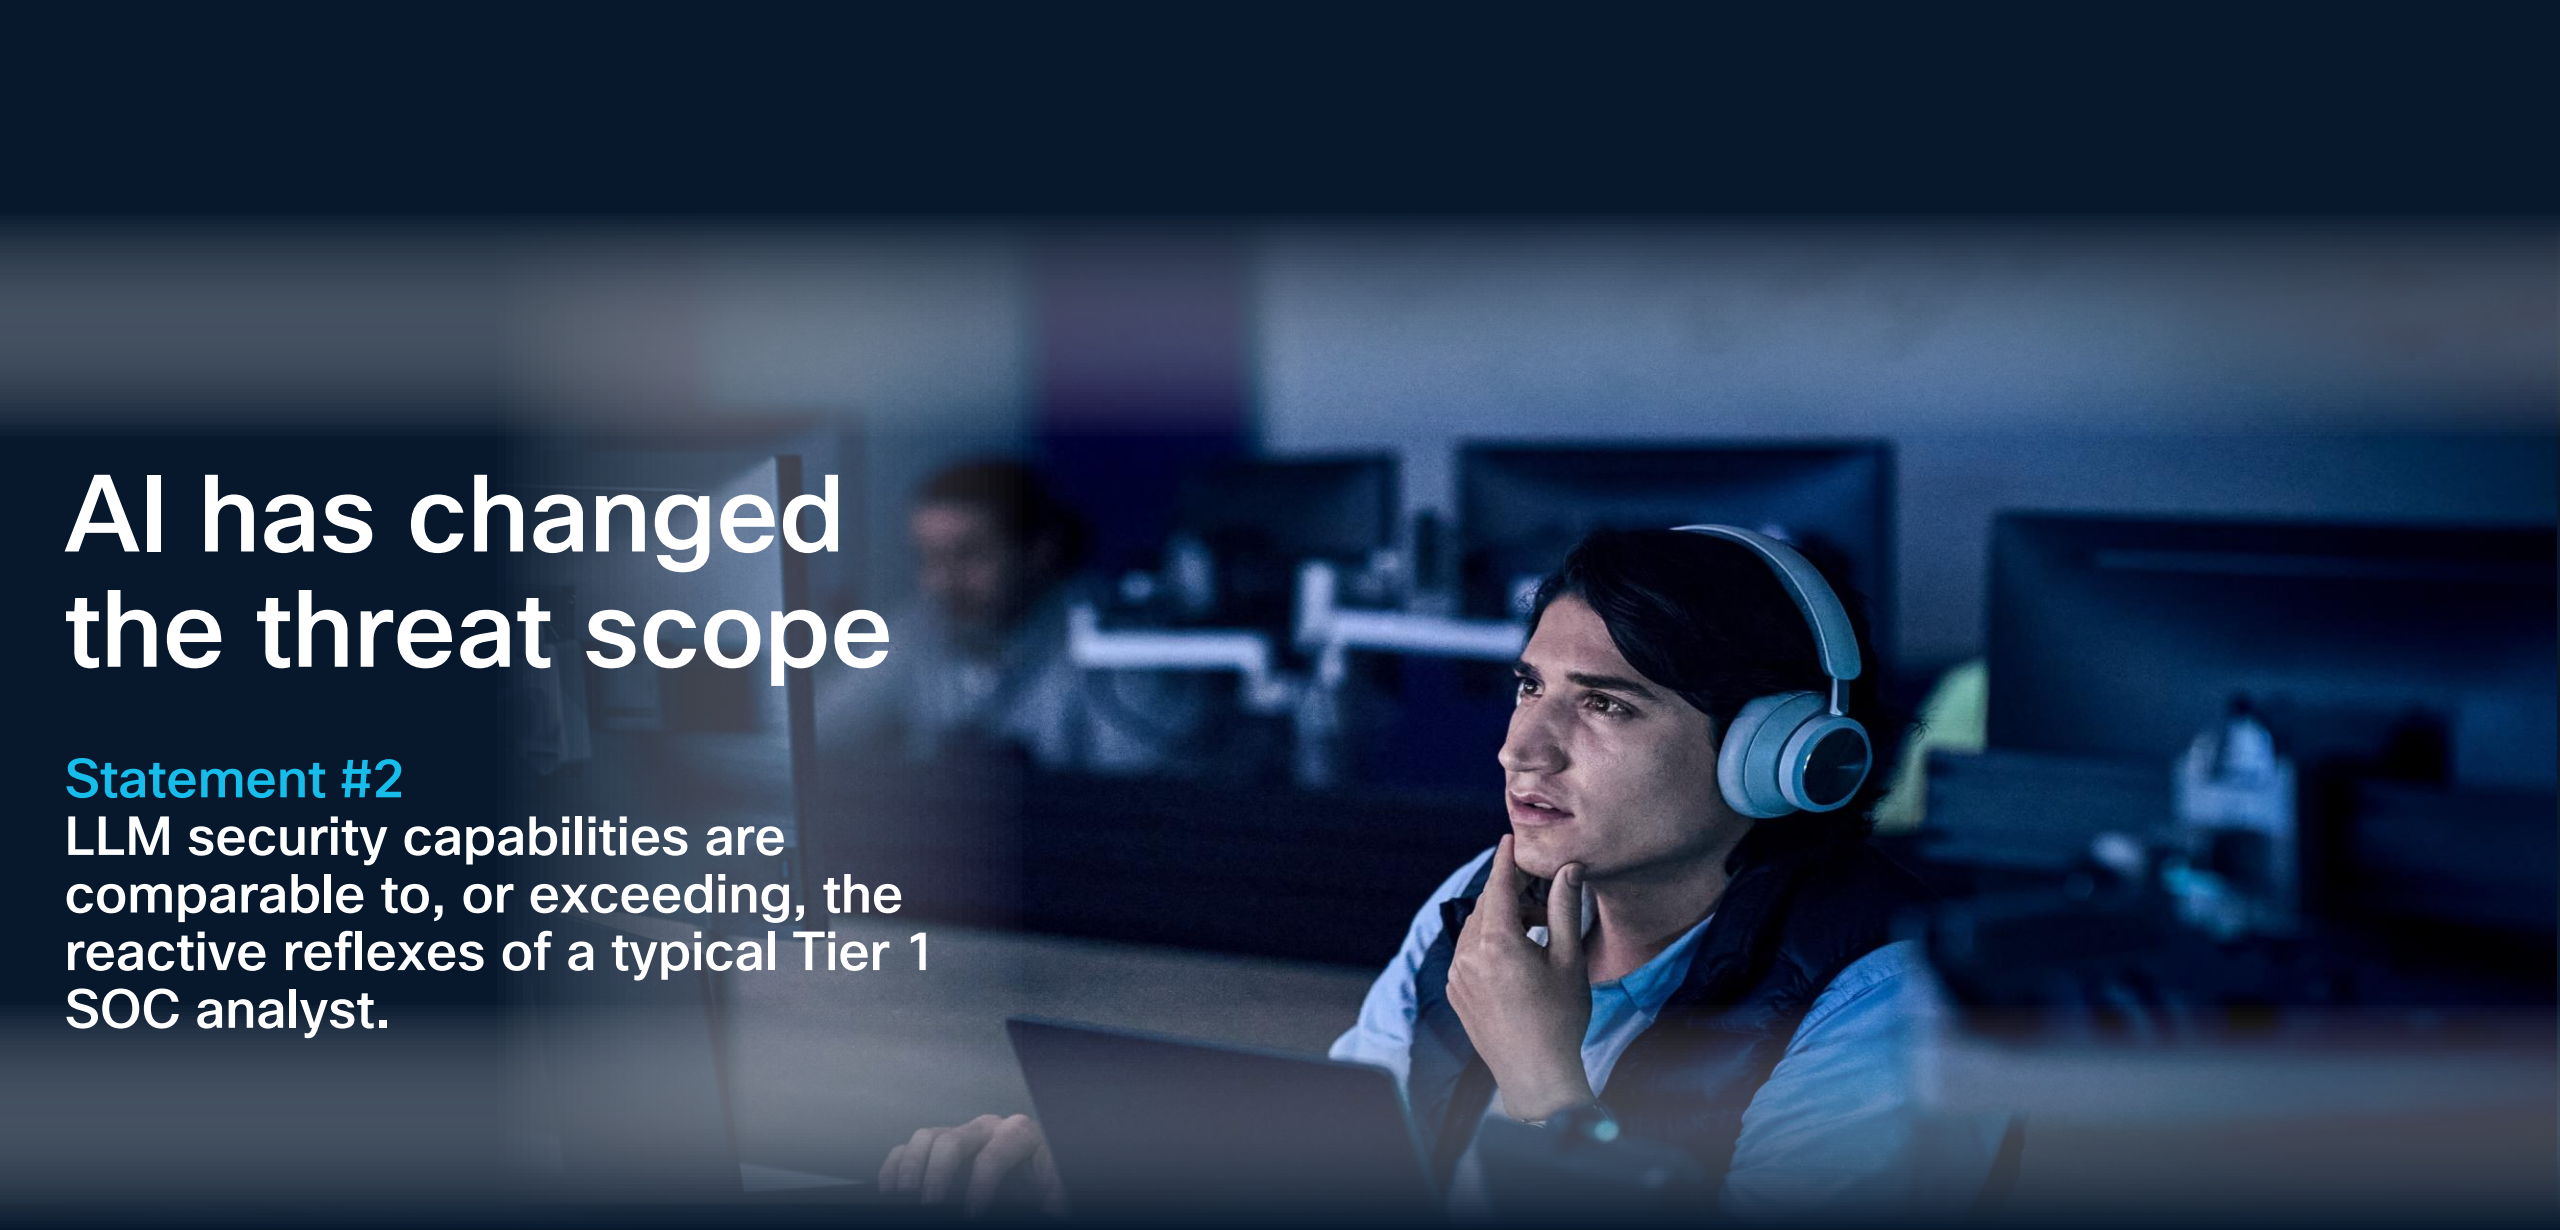

Statement #2

LLM security capabilities are comparable to, or exceeding, the reactive reflexes of a typical Tier 1 SOC analyst.

cisco

AI has changed the threat scope

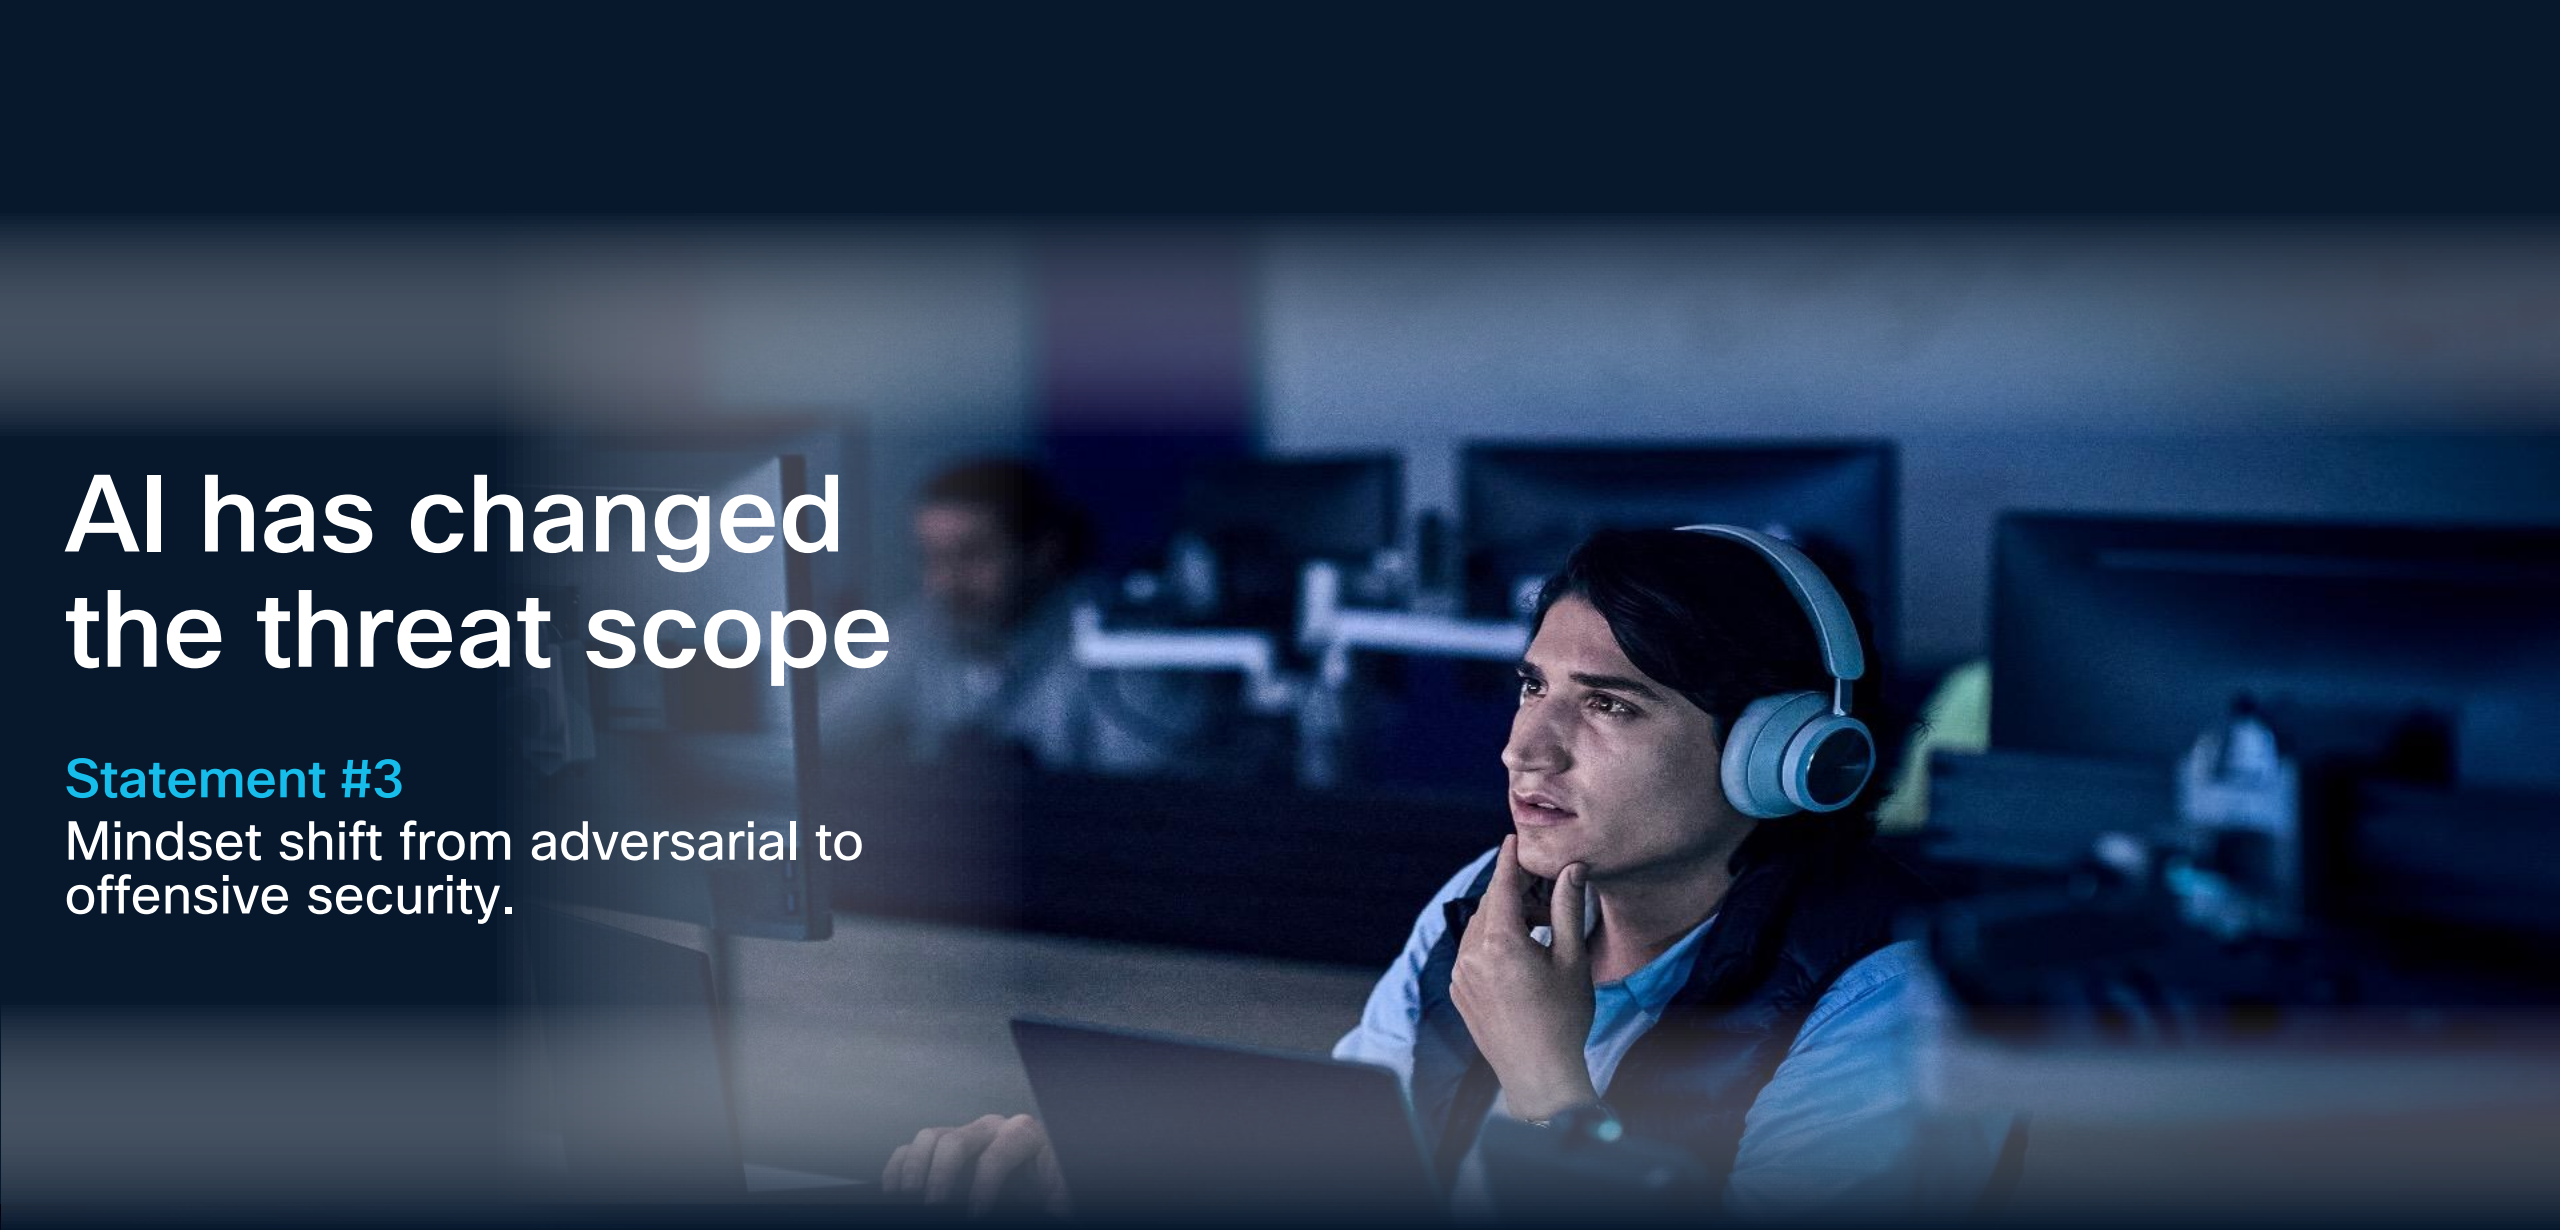

Statement #3

Mindset shift from adversarial to offensive security.

cisco

What AI Can do today?

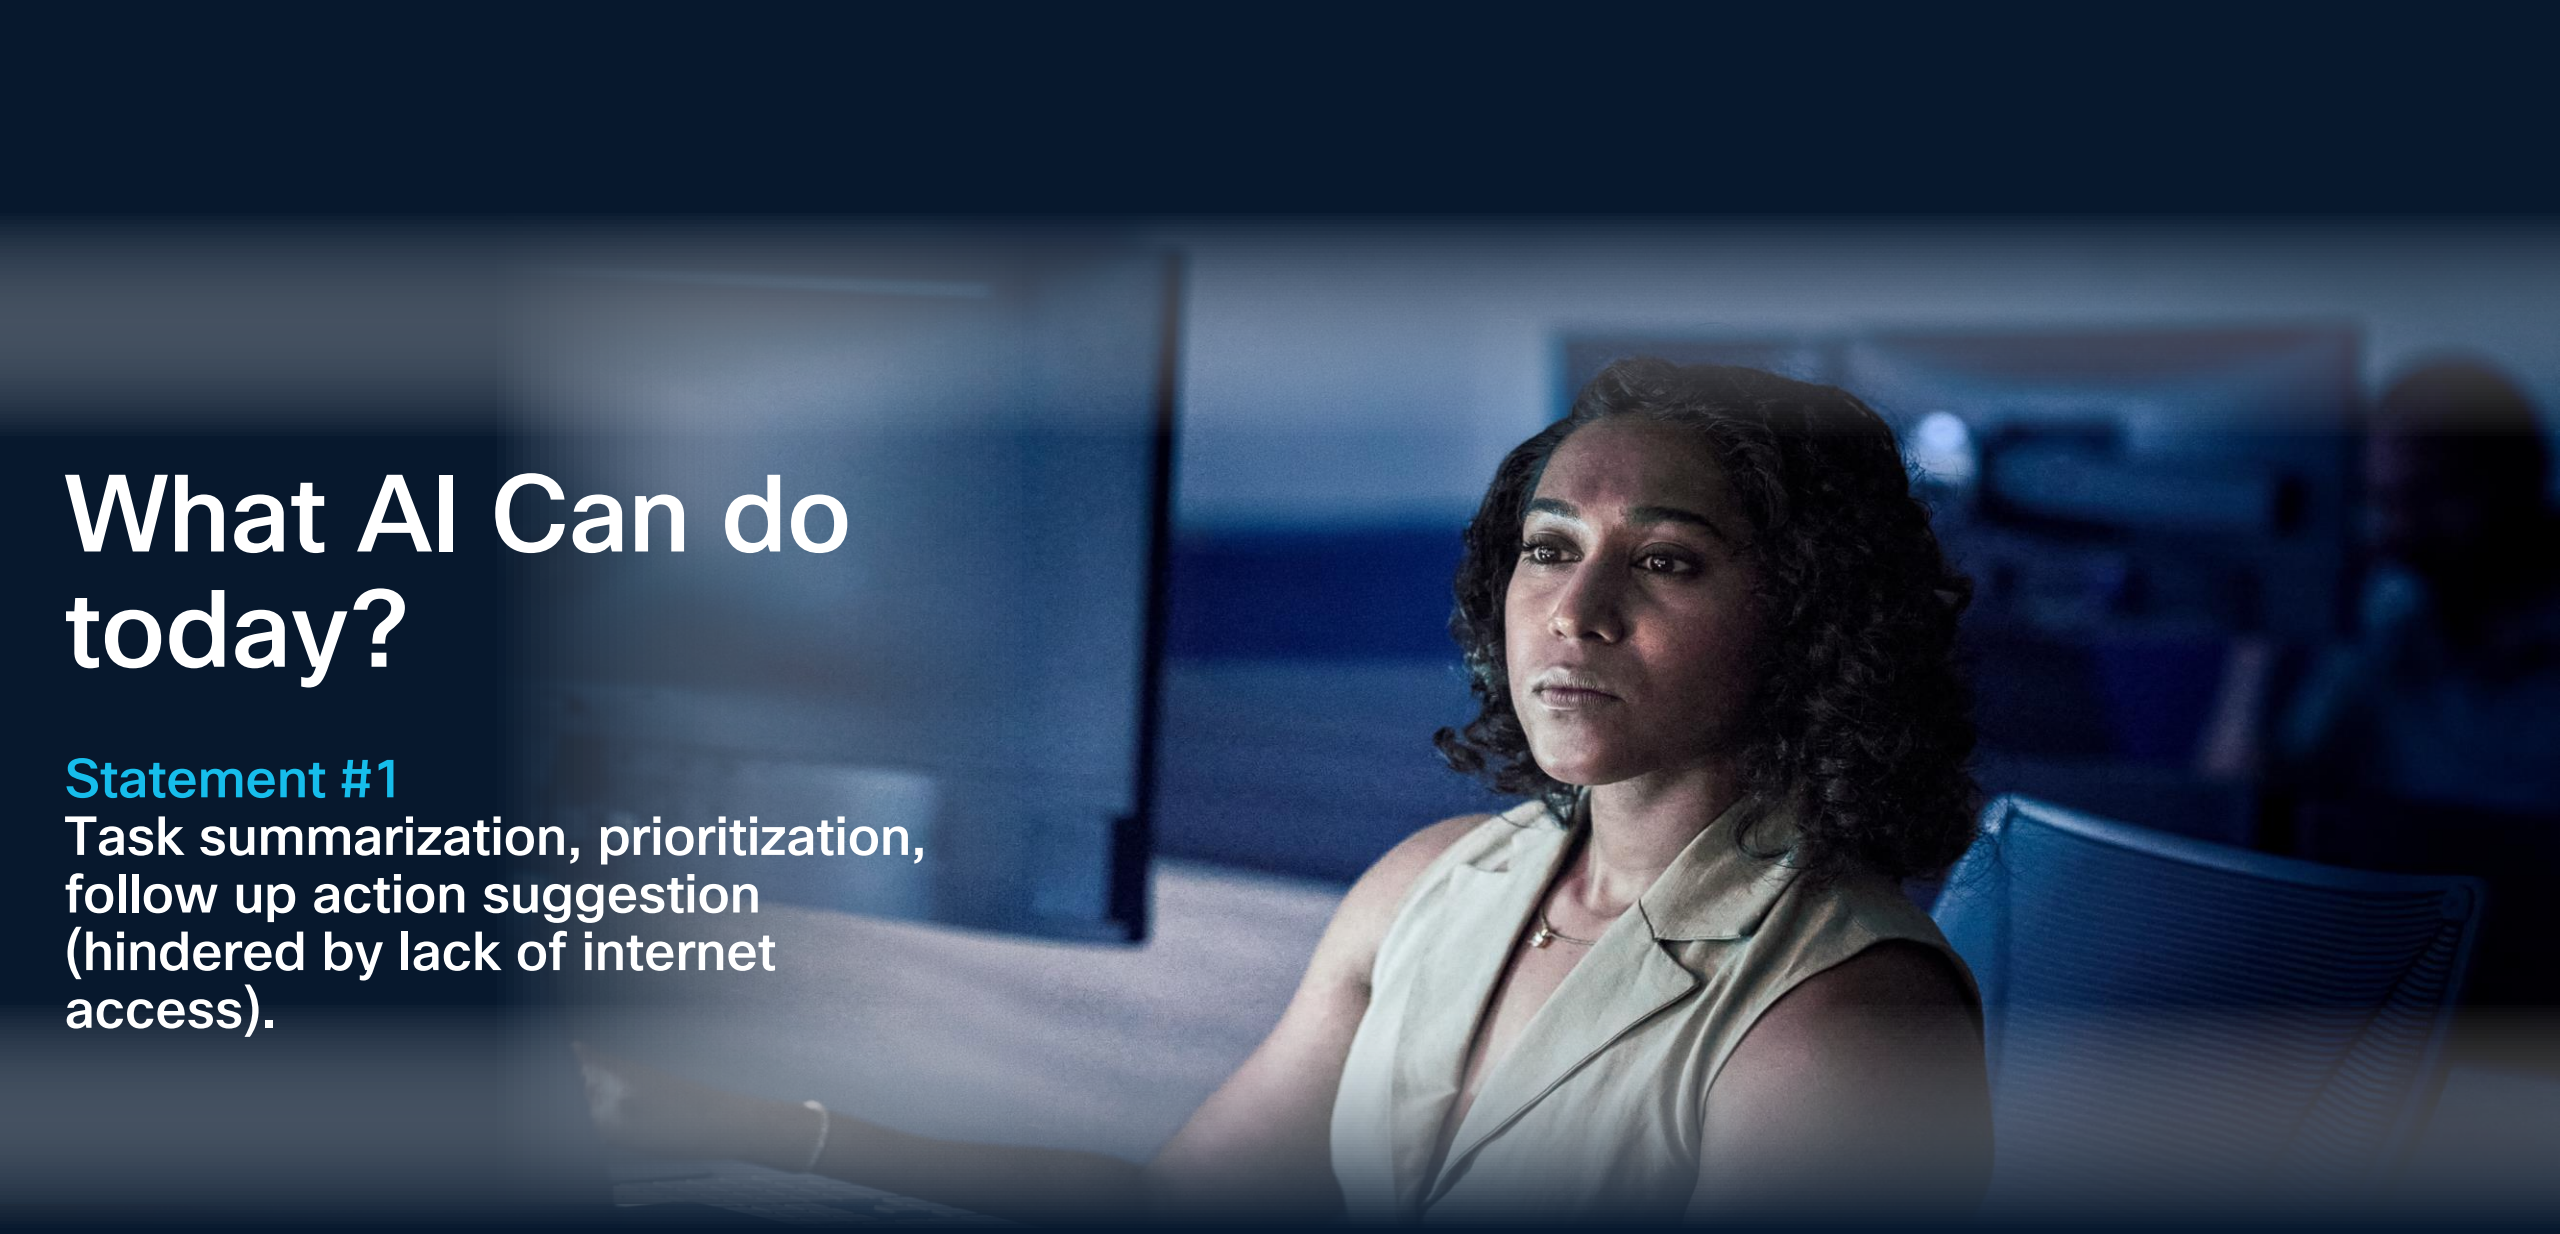

Statement #1

Task summarization, prioritization, follow up action suggestion (hindered by lack of internet access).

cisco

What AI Can do today?

Statement #2

Mislead the user (model sycophancy).

cisco

And SOC?

Statement #1

Generale purpose LLMs make mistakes due to largeness.

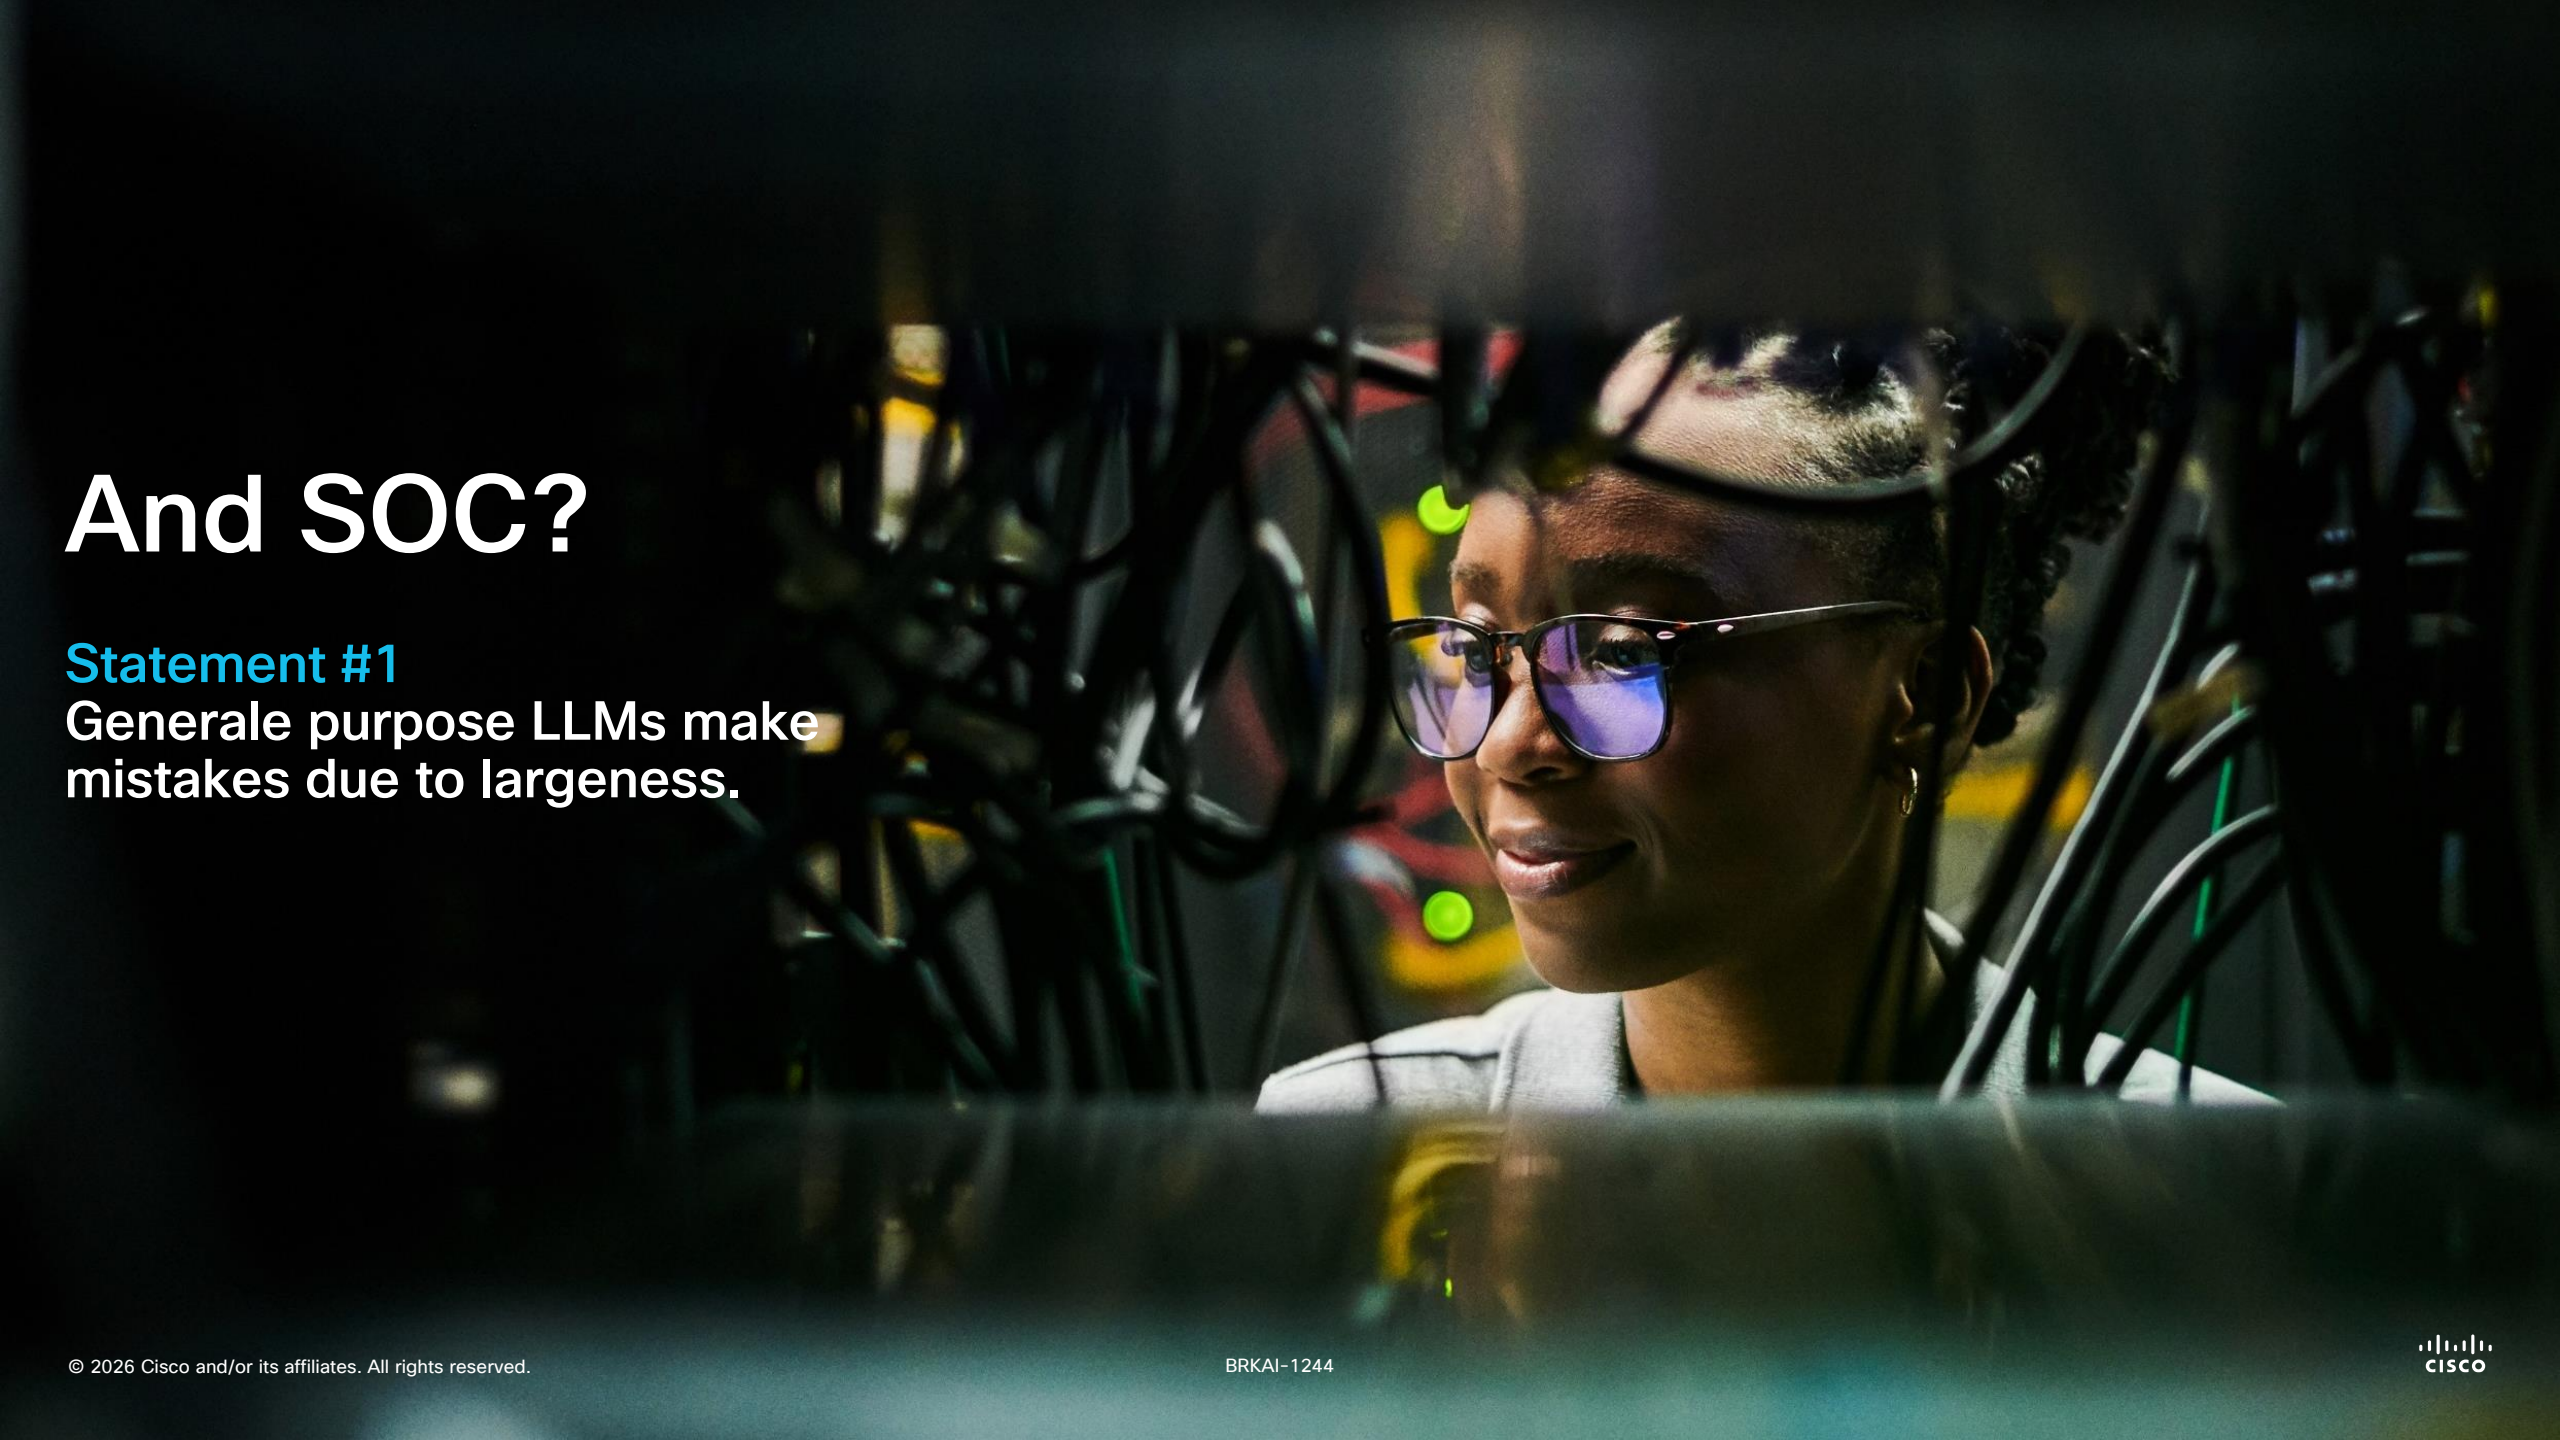

cisco

And SOC?

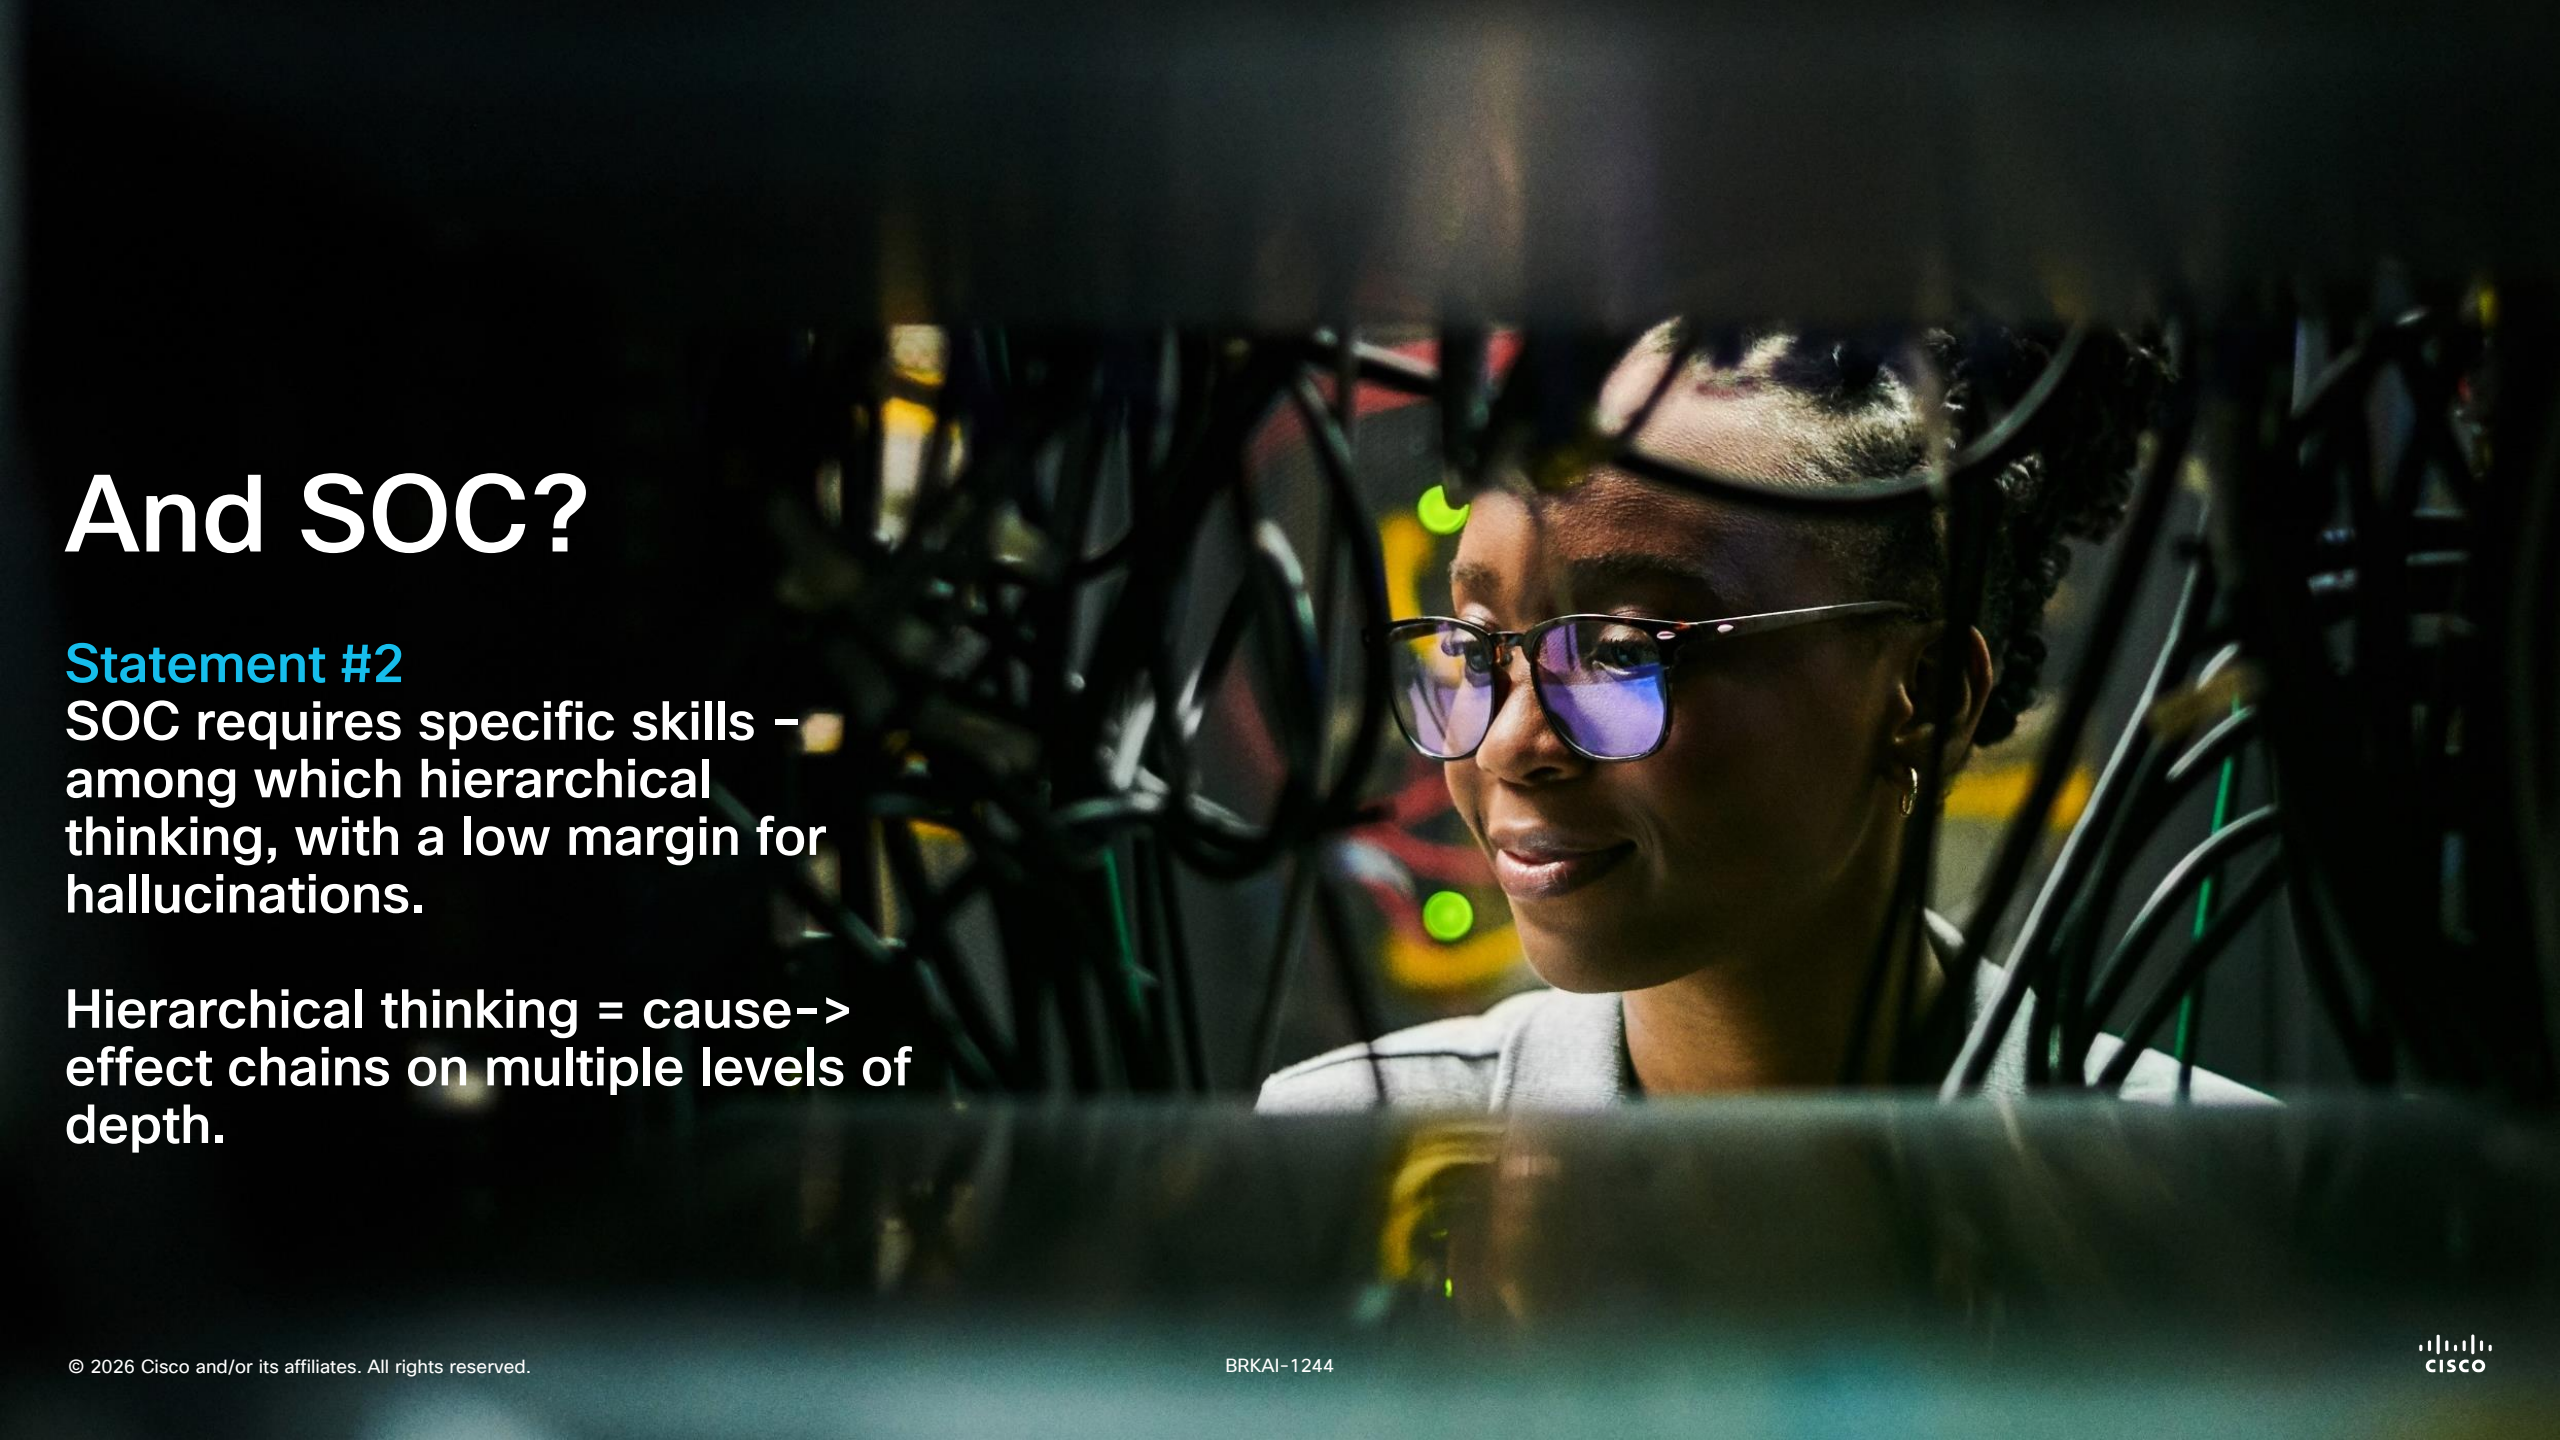

Statement #2

SOC requires specific skills – among which hierarchical thinking, with a low margin for hallucinations.

Hierarchical thinking = cause-> effect chains on multiple levels of depth.

And SOC?

Statement #3

Fine tuning a large language model is less efficient than having more agents cooperating.

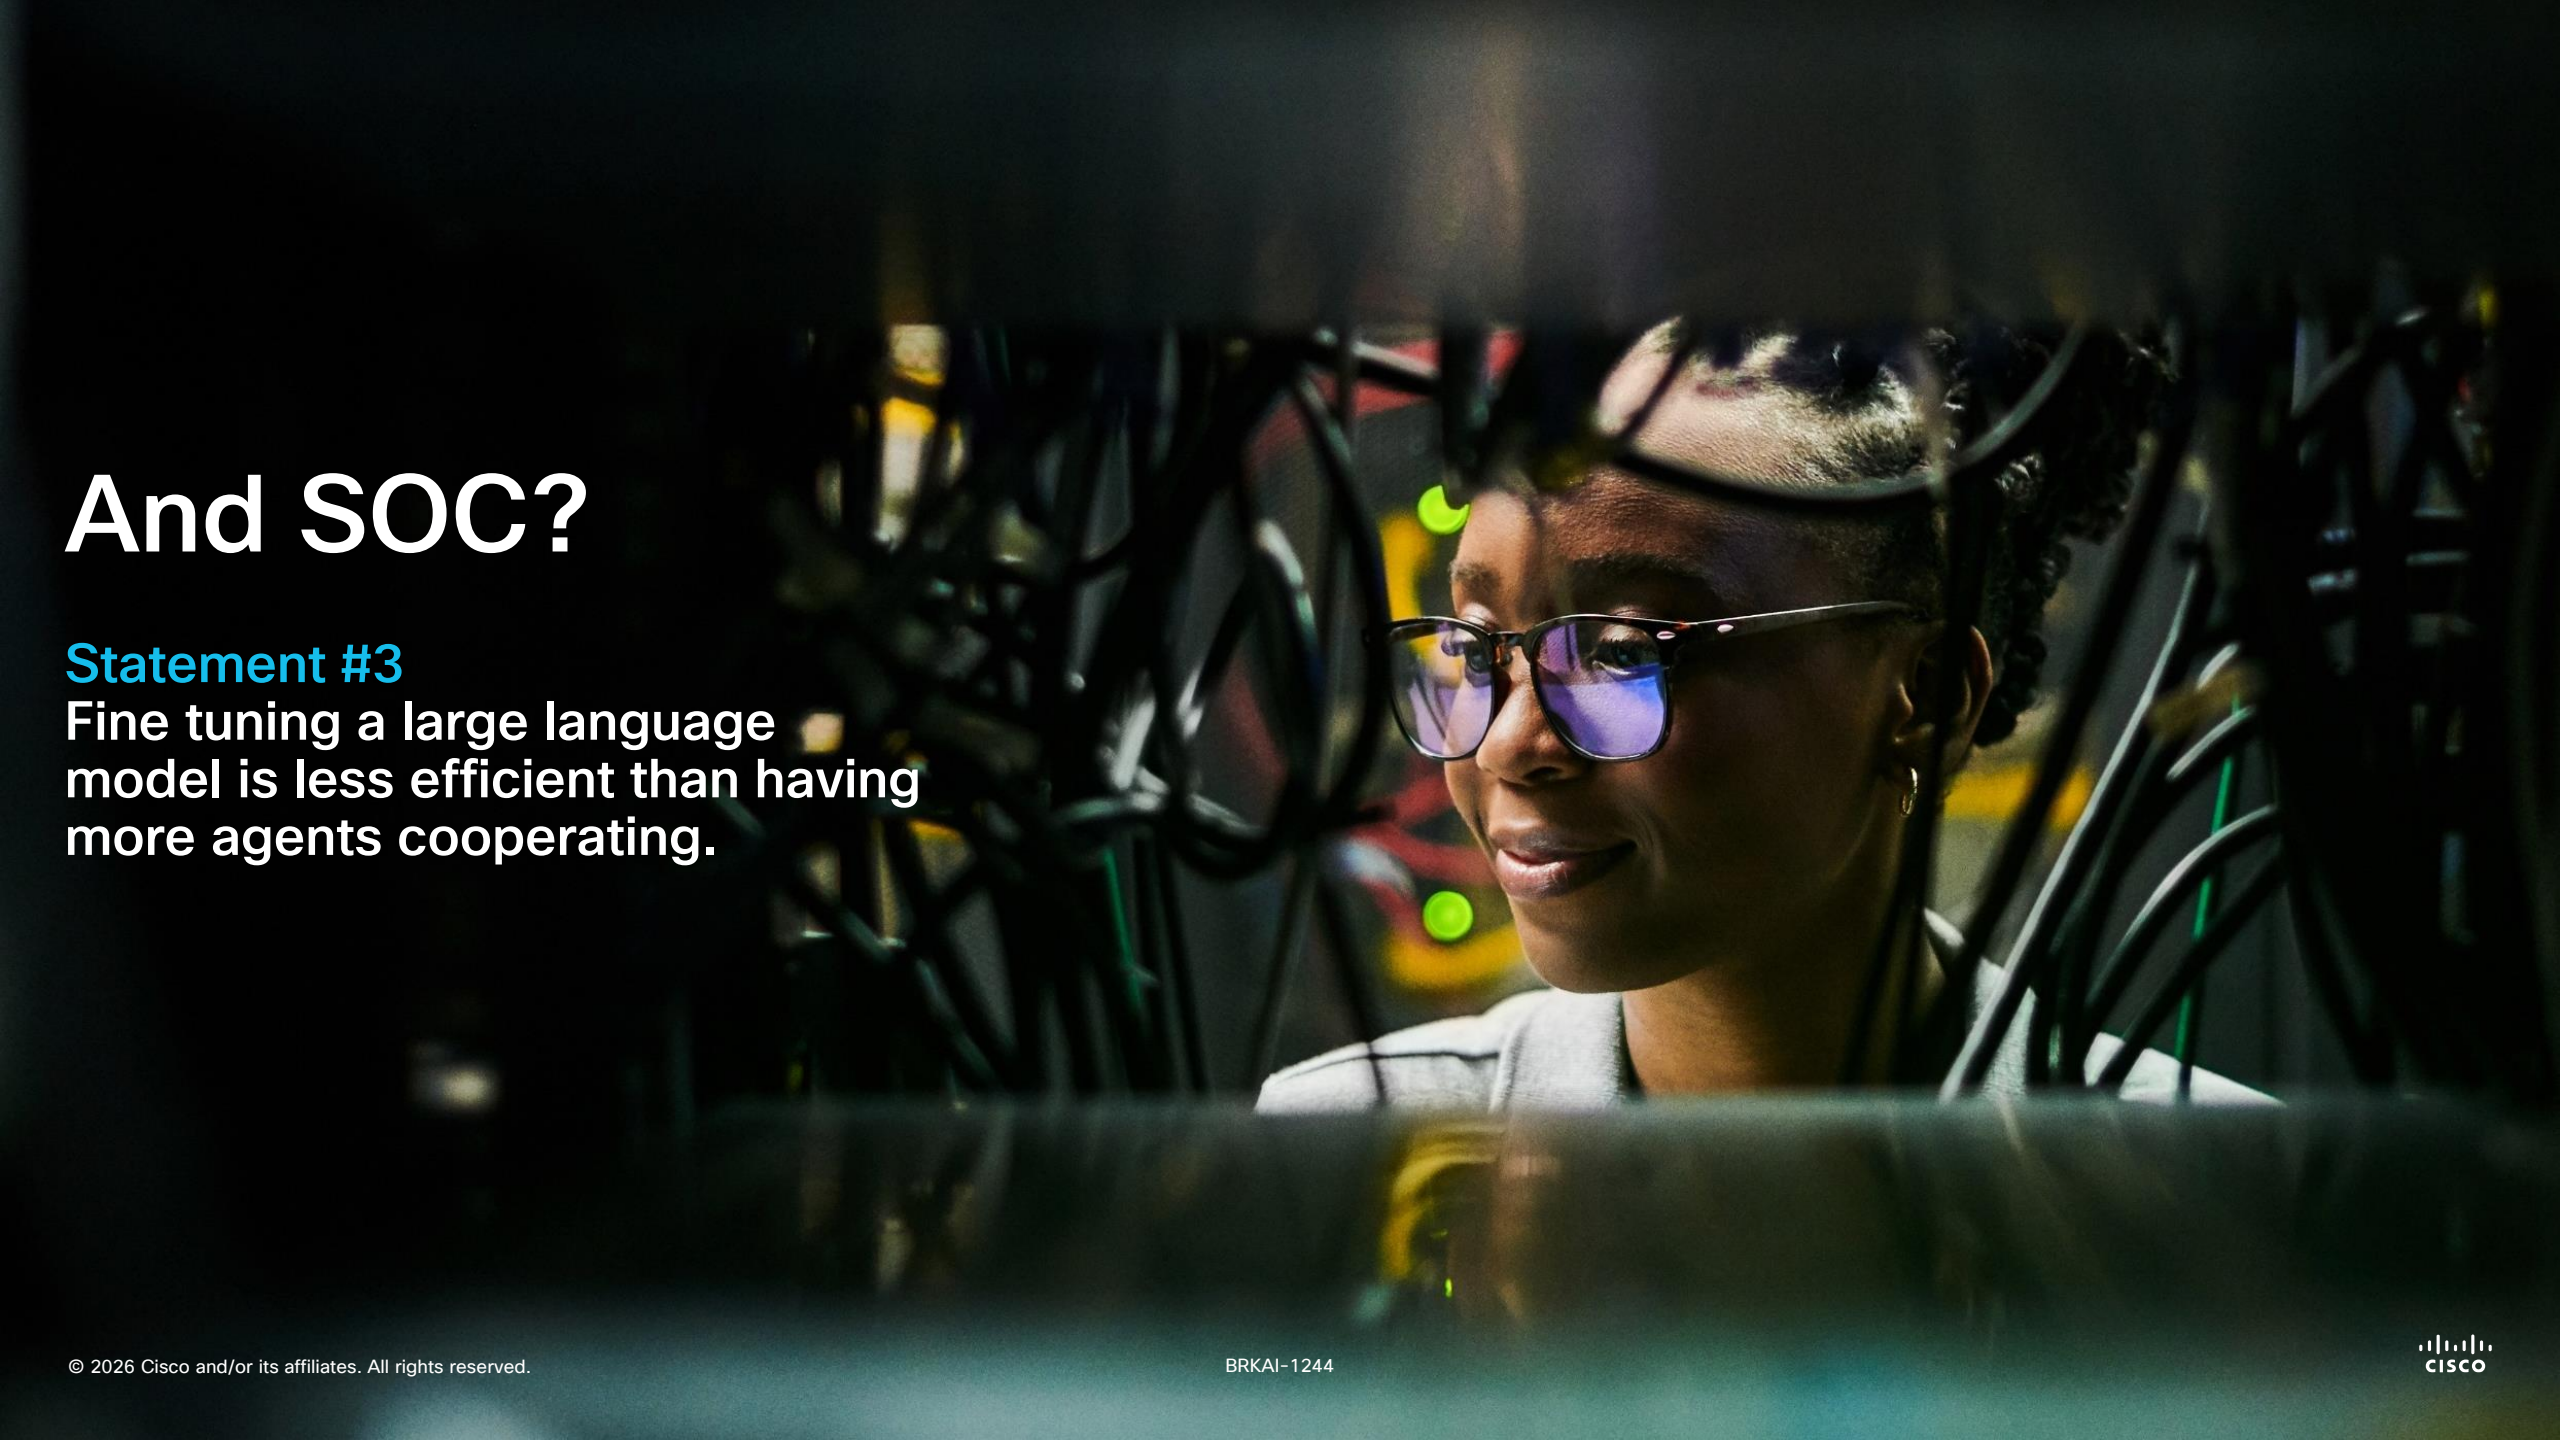

cisco

AI in the SOC: Powerful Tool – Risky if Misplaced

Allow internet browsing:

SLM might leak data or download malware

Retrieval Augmented Generation (RAG)

Expensive, and difficult, you need somewhere to host it.

Vendor lock-in:

Models differ, cloud hosting might be expensive, unpredictable, or outright impossible

Meet the Foundation-Sec-8B Model

"Cisco's acquisition of Robust Intelligence may be one of its best buys in recent years. It provides the automated red-teaming capability in the company's AI Defense platform, and it supports the newly launched Foundation AI open-source reasoning model that is purpose-built for enhancing security applications."

https://www.forbes.com/sites/moorinsights/2025/05/07/rsa-conference-2025-highlights-insights-and-companies-to-watch/

Foundation AI Background

- Spun out from the **Robust Intelligence** team
- Established in **March 2025** with a newly defined charter
- Core mission: build bleeding edge AI for cybersecurity and cybersecurity for AI
- Highly experienced team with decades of combined AI and security expertise

Foundation AI OSS Security Model Background

Open-weight, **8B parameter** model, Cisco's first domain-specific LLM for cybersecurity

Pre-trained on **Llama-3.1 8B**

Designed from the ground up for security workflows

Trained on a curated corpus of **CVEs**, **CWEs**, **MITRE ATT&CK**, threat-intel reports, red-team playbooks, tooling docs and compliance standards, delivering analyst-ready context with fewer hallucination

Development Highlights

Building infrastructure for model pre-training, finetuning, and model evaluation

Building a novel data collection system

- Web scrapers visited over 160 million domains and 60 large security specific sources, yielding over 4 TiB of raw unfiltered data.
- Raw data is then passed through language, quality, and relevancy filters before being deduplicated and undergoing PII removal.
- Pre-training on 7B high-quality cybersecurity tokens.

Handpicked special sources from trusted security and research domains:

- Cyber Threat Intel: MITRE ATT&CK, CVE, CWE, NVD, OWASP, GitHub Advisories
- Incident Response: NIST CSRC guides, SOC/IR playbooks, RFCs, PRs, postmortems
- Cloud & Tooling Docs: GCP, Azure, Kubernetes
- Research & Textbooks: ArXiv, Open Textbooks
- Standards & Policy: HIPAA, SOC2, public regulations

Cisco Foundation-AI: Built in 3 Forms, Designed for Every Need

Foundation-Sec-8B

Foundation-Sec-8B Instruct

Foundation-Sec-8B Reasoning

cisco

Understanding Foundation AI Models: Instruct and Reasoning Overview

Models Everywhere

Foundation Model

Large, pre-trained AI on diverse data;

Flexible for many tasks.

Instruct Model

Fine-tunes foundation models to follow user instructions better, improving task-specific responses.

Reasoning Model

Enables complex, multi-step logical thinking for advanced analysis and decision-making in specialized domains.

Ok, where to get the Foundation AI Sec-8b models?

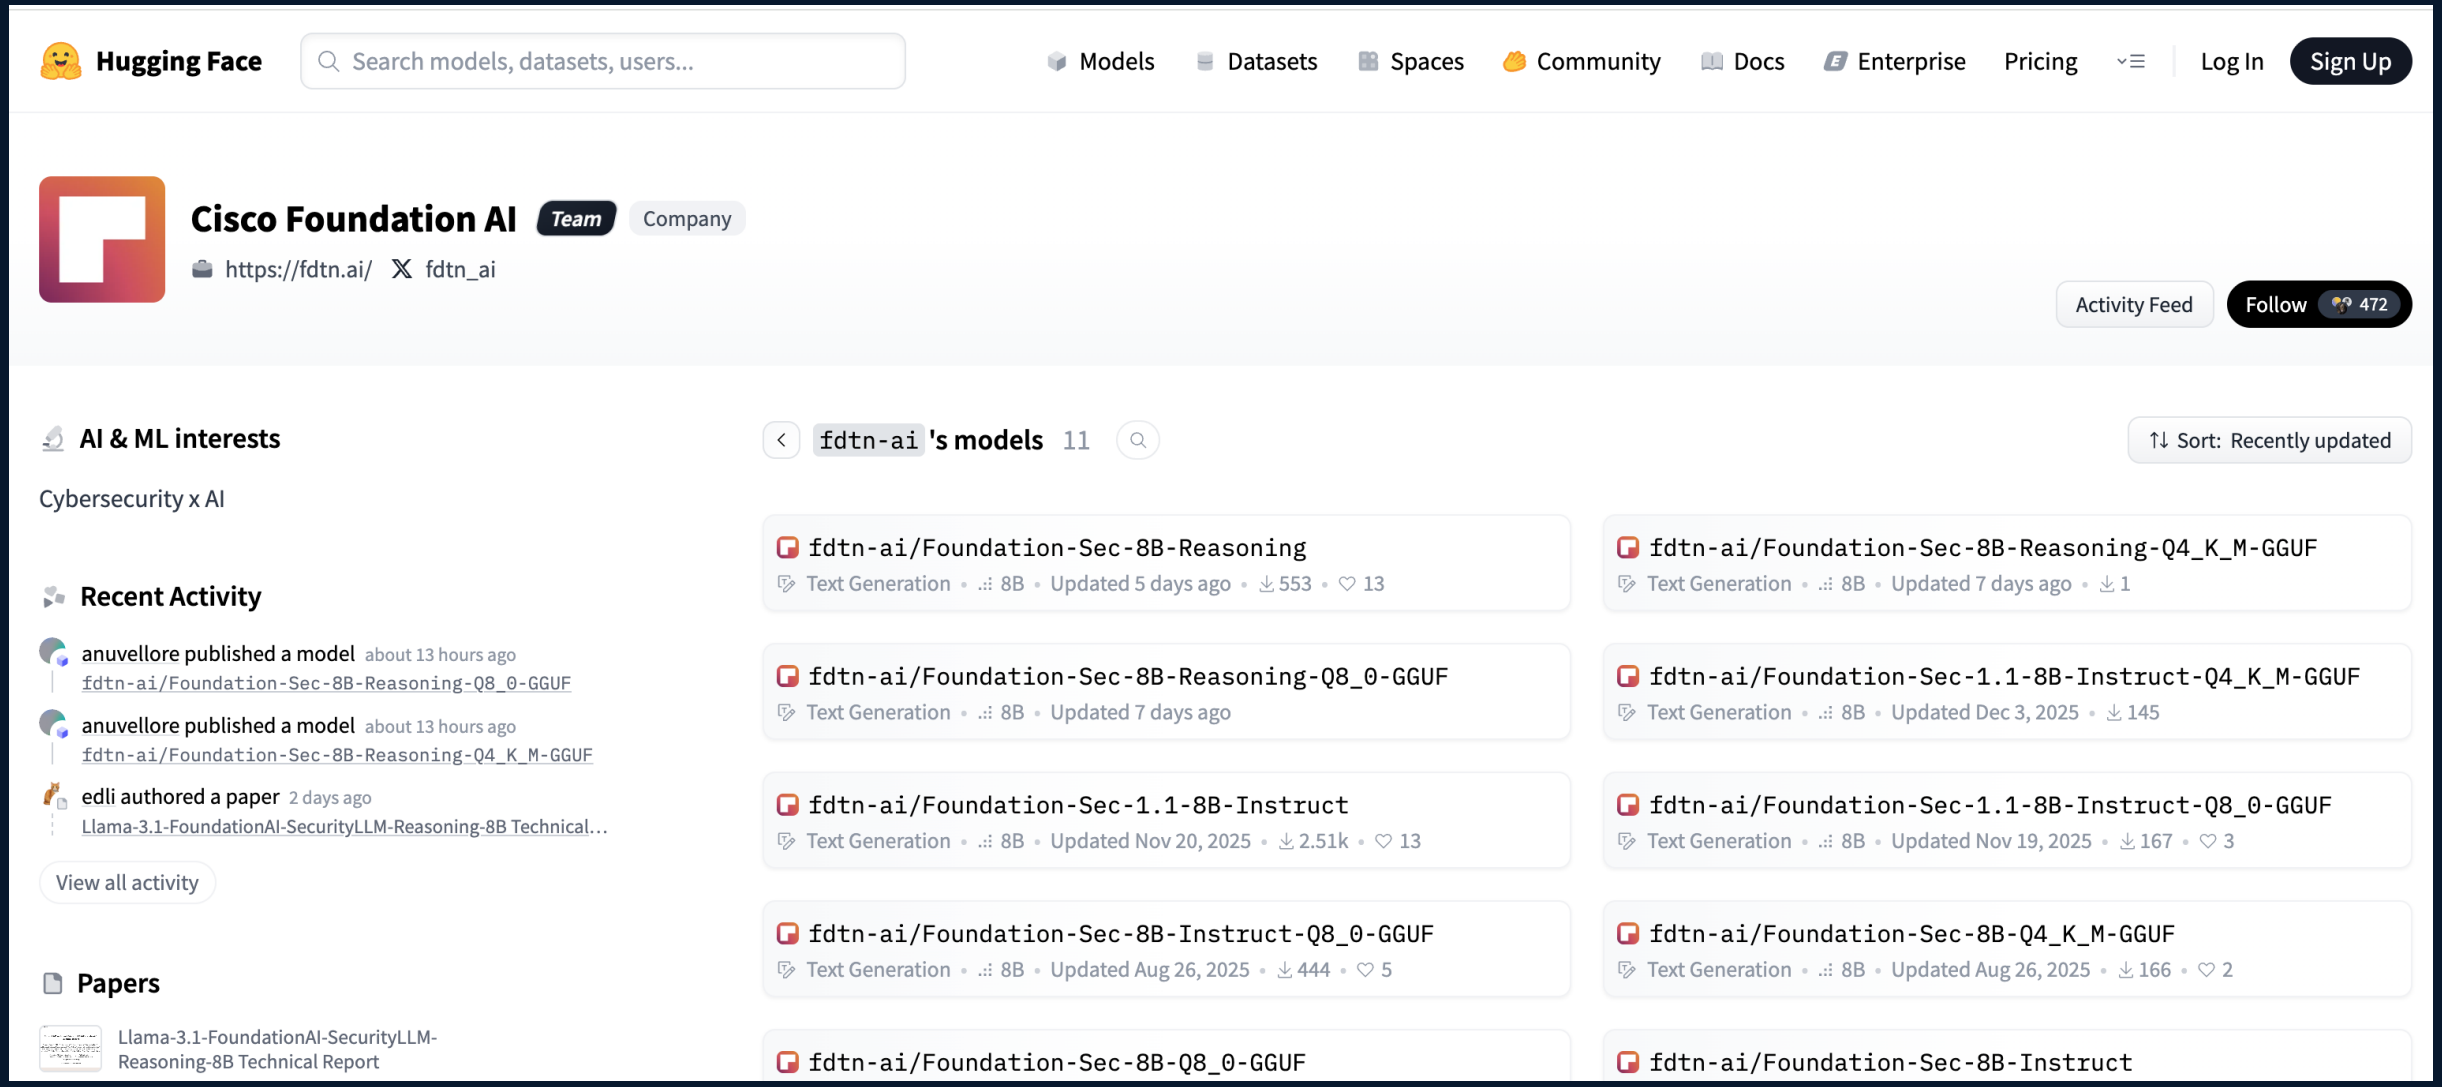

Sec-8B Model Use Cases

Foundation security models enable the SOC

SOC Teams are overwhelmed by:

- False Positives
- Manual Repetitive Triage
- Fragmented Tools

Model is ready to embed across the **SOC lifecycle**:

- Advance Threat Detection & Investigation
- Adversary Behavior & Threat Intelligence
- Security Posture & Vulnerability Management
- Proactive Defense & Engineering Enablement

Model is being **adopted** in **Pilots** across Cisco:

- Splunk
- Cisco CSIRT
- Cisco XDR

What’s difference with other models?

You are a senior SOC analyst assisting with incident triage.Your task is to read the incident metadata and alert logs, and provide a clear summary of what occurred. Instructions:
- Highlight the sequence of events (inferred from timestamps).
- Think deeply about cause and effect and how artifacts relate to one another
- Mention key attack techniques used (if inferable from logs)
- Describe how the attack began and progressed
- Use clear and concise language appropriate for L1/L2 analysts

incident_id: INC-1024,
type: Unauthorized Access,
severity: High
timestamp: 2025-04-09T10:30:00Z
2025-04-09 10:00:23 - Alert: 5 failed login attempts for user 'alice' on host 'WS123'
2025-04-09 10:05:10 - Alert: Suspicious PowerShell execution on 'WS123' by 'alice' (malicious script blocked)
2025-04-09 10:10:45 - Alert: Process dumping LSASS memory on 'WS123' (possible credential theft)
2025-04-09 10:15:00 - Alert: Successful login of 'alice' to server 'DC1' from host 'WS123'
Summarize what happened in this incident in a few sentences

What’s difference with other models?

mistral

Temporary Chat

Here's a summary of the incident **INC-1024**:

Sequence of Events:

1. **2025-04-09 10:00:23**: Suspect made 5 failed login attempts for user 'alice' on host 'WS123'.
2. **2025-04-09 10:05:10**:With failed login attempts complete, the suspect executed a malicious powershell script.
3. **2025-04-09 10:10:45**: The suspect attempted to dump LSASS memory on 'WS123', likely trying to steal credentials.
4. **2025-04-09 10:15:00**: The suspect successfully logged into the domain controller 'DC1' using 'alice's credentials'.

Key Attack Techniques Used:

- Credential dumping. The attack started with brute force login attempts. After failing, the attacker tried to execute a malware but caused it to be blocked.
- Credential Access: Attempted to dump LSASS memory to extract credentials which may have worked for the moment since they successfully used the account.

Incident Progression:

- The attacker began with a brute-force attempt to gain access to the account. They then executed a suspicious PowerShell script, likely to bypass security measures. After likely failing to get credentials using brute force the attacker used a credential access method to get credentials. Using those, the attacker successfully accessed the domain controller 'DC1'. The potential motive: Domain Admin access. Patching the access any server using Alice's account would be essential.

Suspicious behaviors that may have indicated an attacker:

- 5 failed login attempts for the same user
- Suspicious PowerShell execution in a timeframe to security logs.
- Using a legitimate account to attempt additional escalation.
- Unusual remote logins for a user immediately after credible access.

Likely the attacker wanted to get an initial foothold into the system and quickly tried credential-based exploits to elevate their access to administrative-level controls in a timeframe of about 15 minutes

What’s difference with other models?

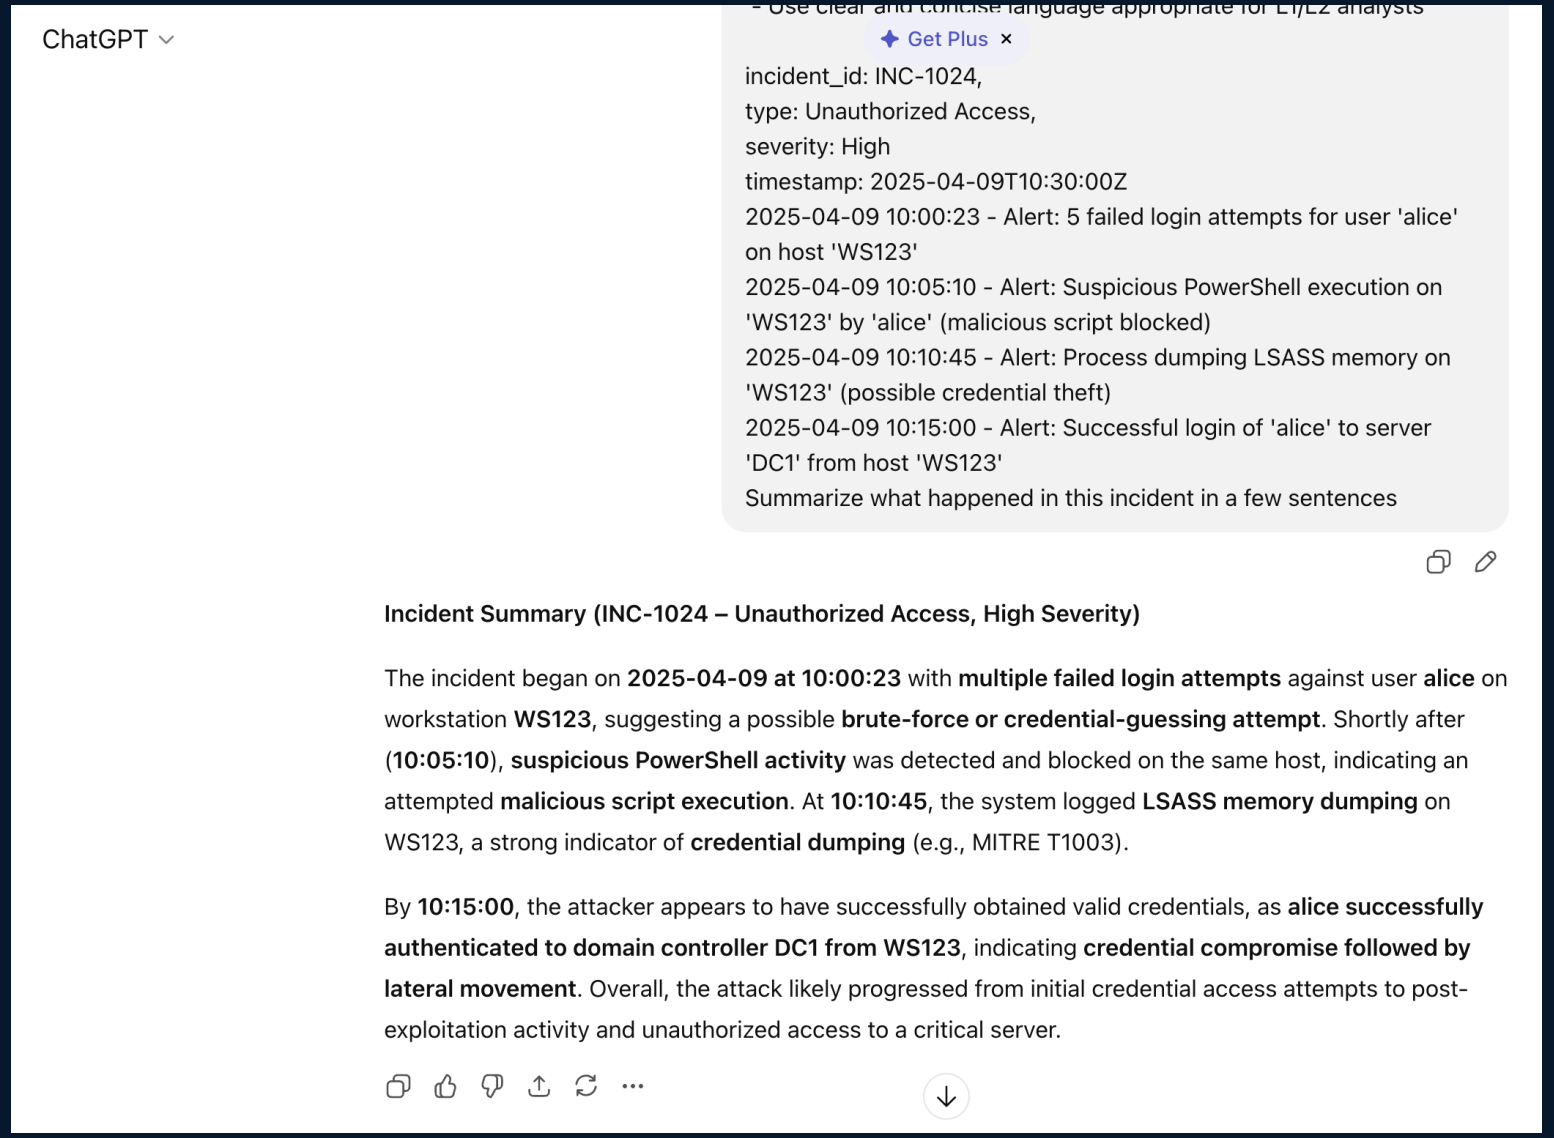

What’s difference with other models?

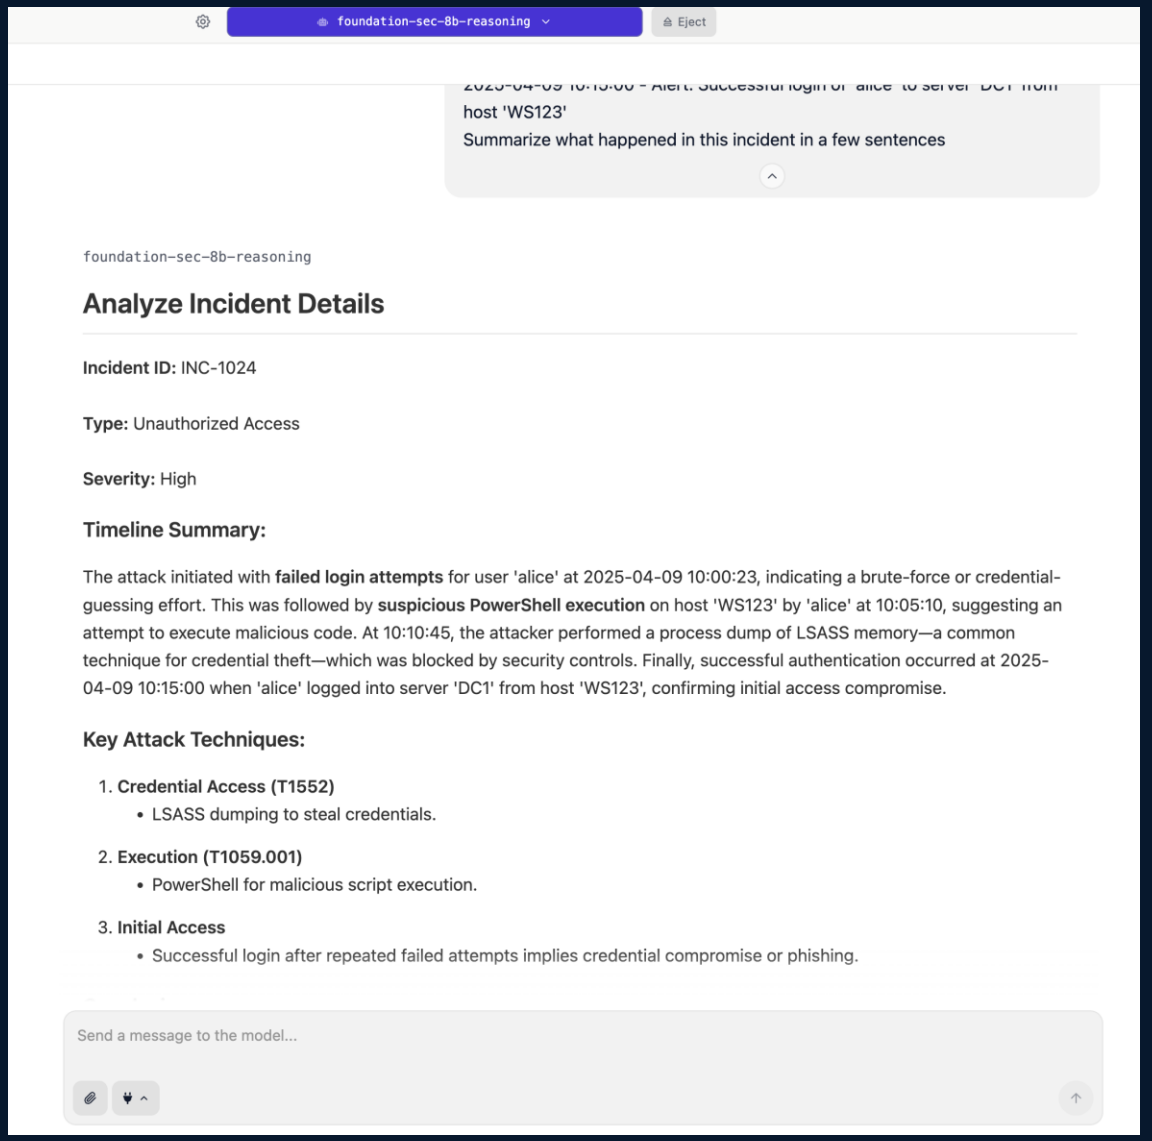

Meet the Cisco AI POD

What’s Needed to Run Foundation-AI Sec-8b LLM?

Sec-8B can run on laptop, but we want centralized access and carrier grade reliability

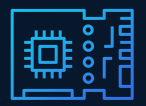

GPU Acceleration

Enables simultaneous inferencing for many users through optimized GPU partitioning and orchestration.

Infrastructure Essential

Modular systems to enable independent scaling of CPUs, GPUs, memory, and storage.

Software and Management

Cloud-based approach that manages multi-tenant AI Workloads with integrated management tools.

Security and Compliance

Compliance frameworks and observability tools ensure secure, resilient operations across all users.

Robust, flexible infrastructure to meet your needs

Data center building blocks

Nexus

Unified Computing System

cisco

Robust, flexible infrastructure to meet your needs

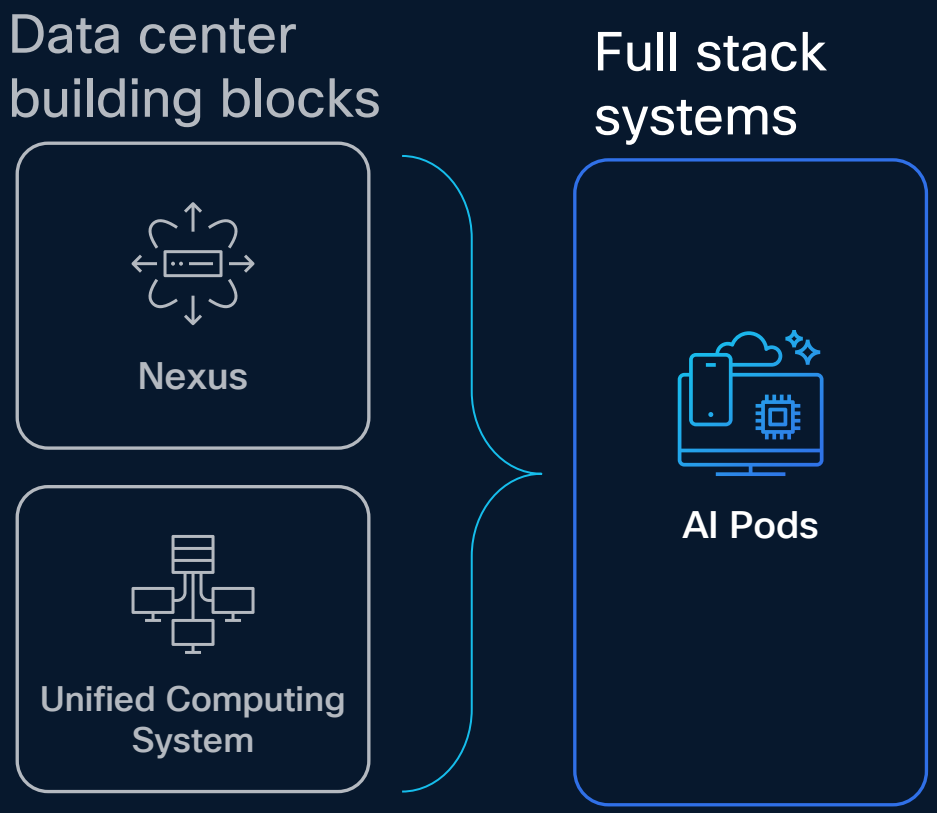

Robust, flexible infrastructure to meet your needs

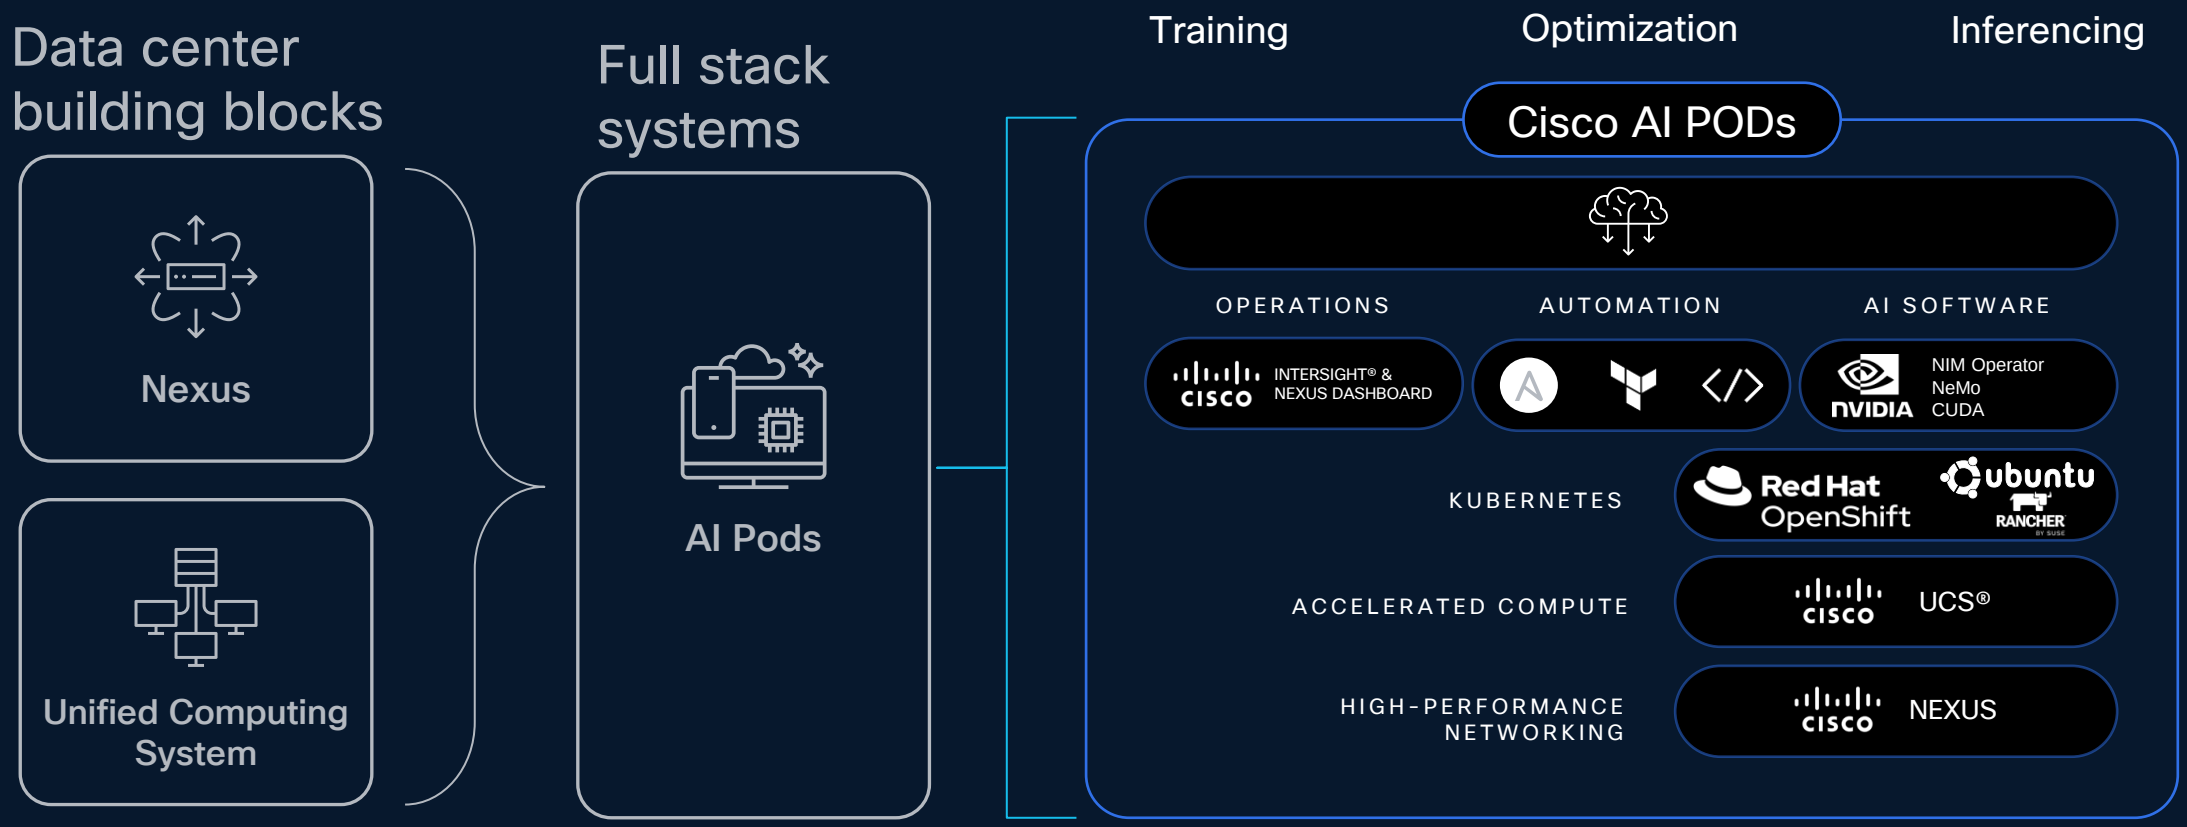

Cisco AI Pod Full Stack System

What is an AI POD?

Modular components for flexible adoption

Deploy AI with confidence: CVD

Sizing tools for predictability

Fully stack by Cisco, including 3rd Party

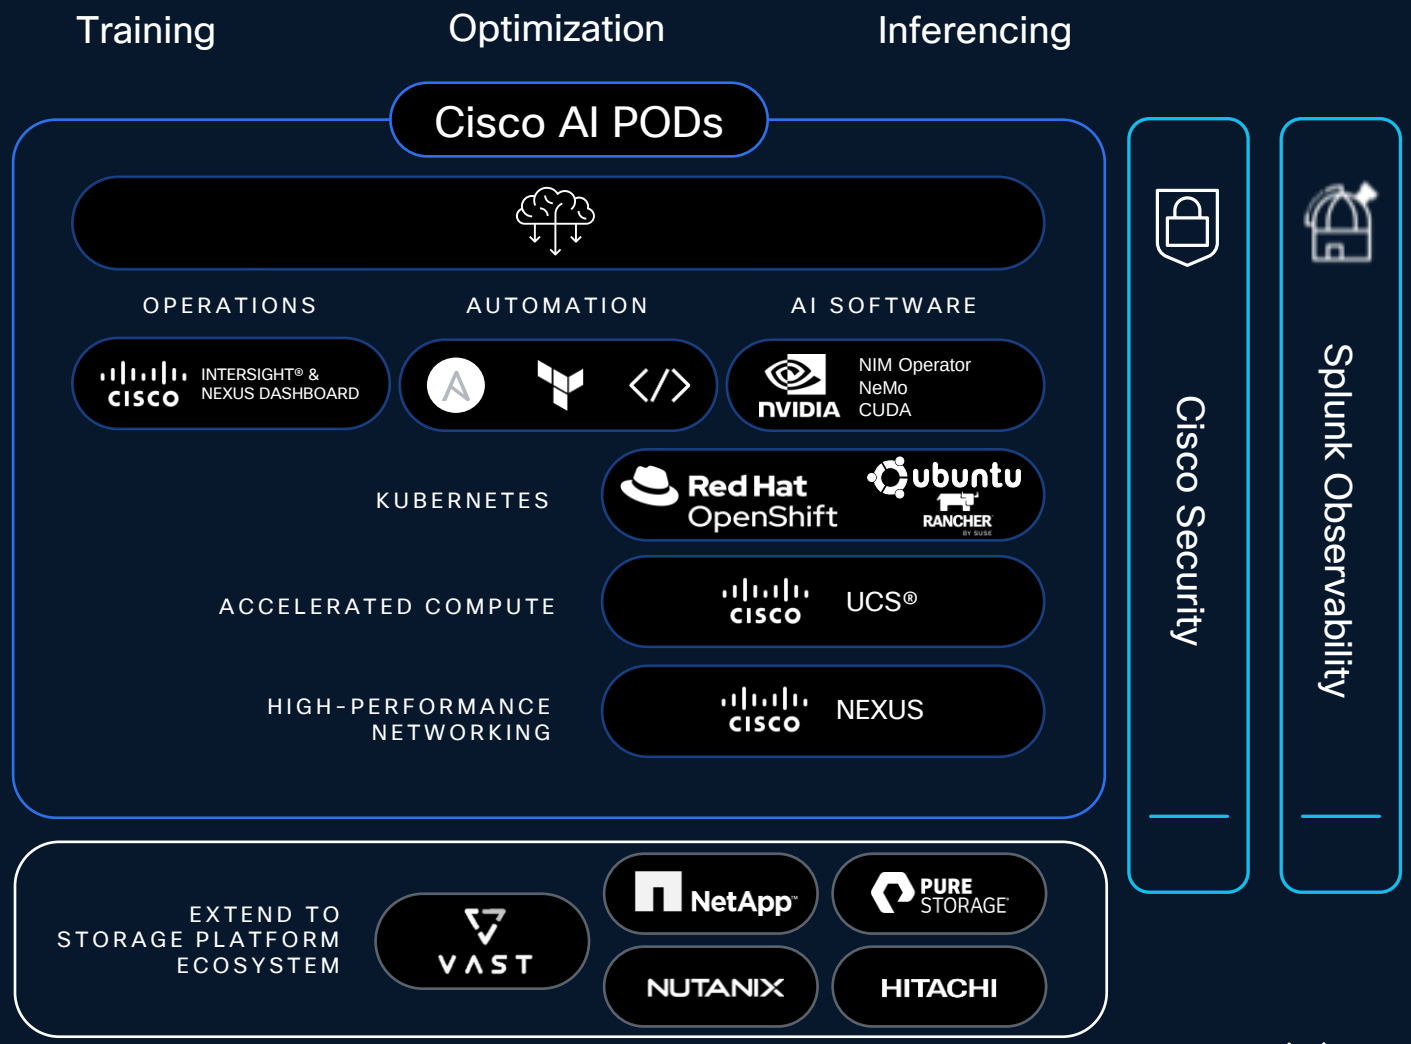

LLM for SOC: Sample Executive Architecture

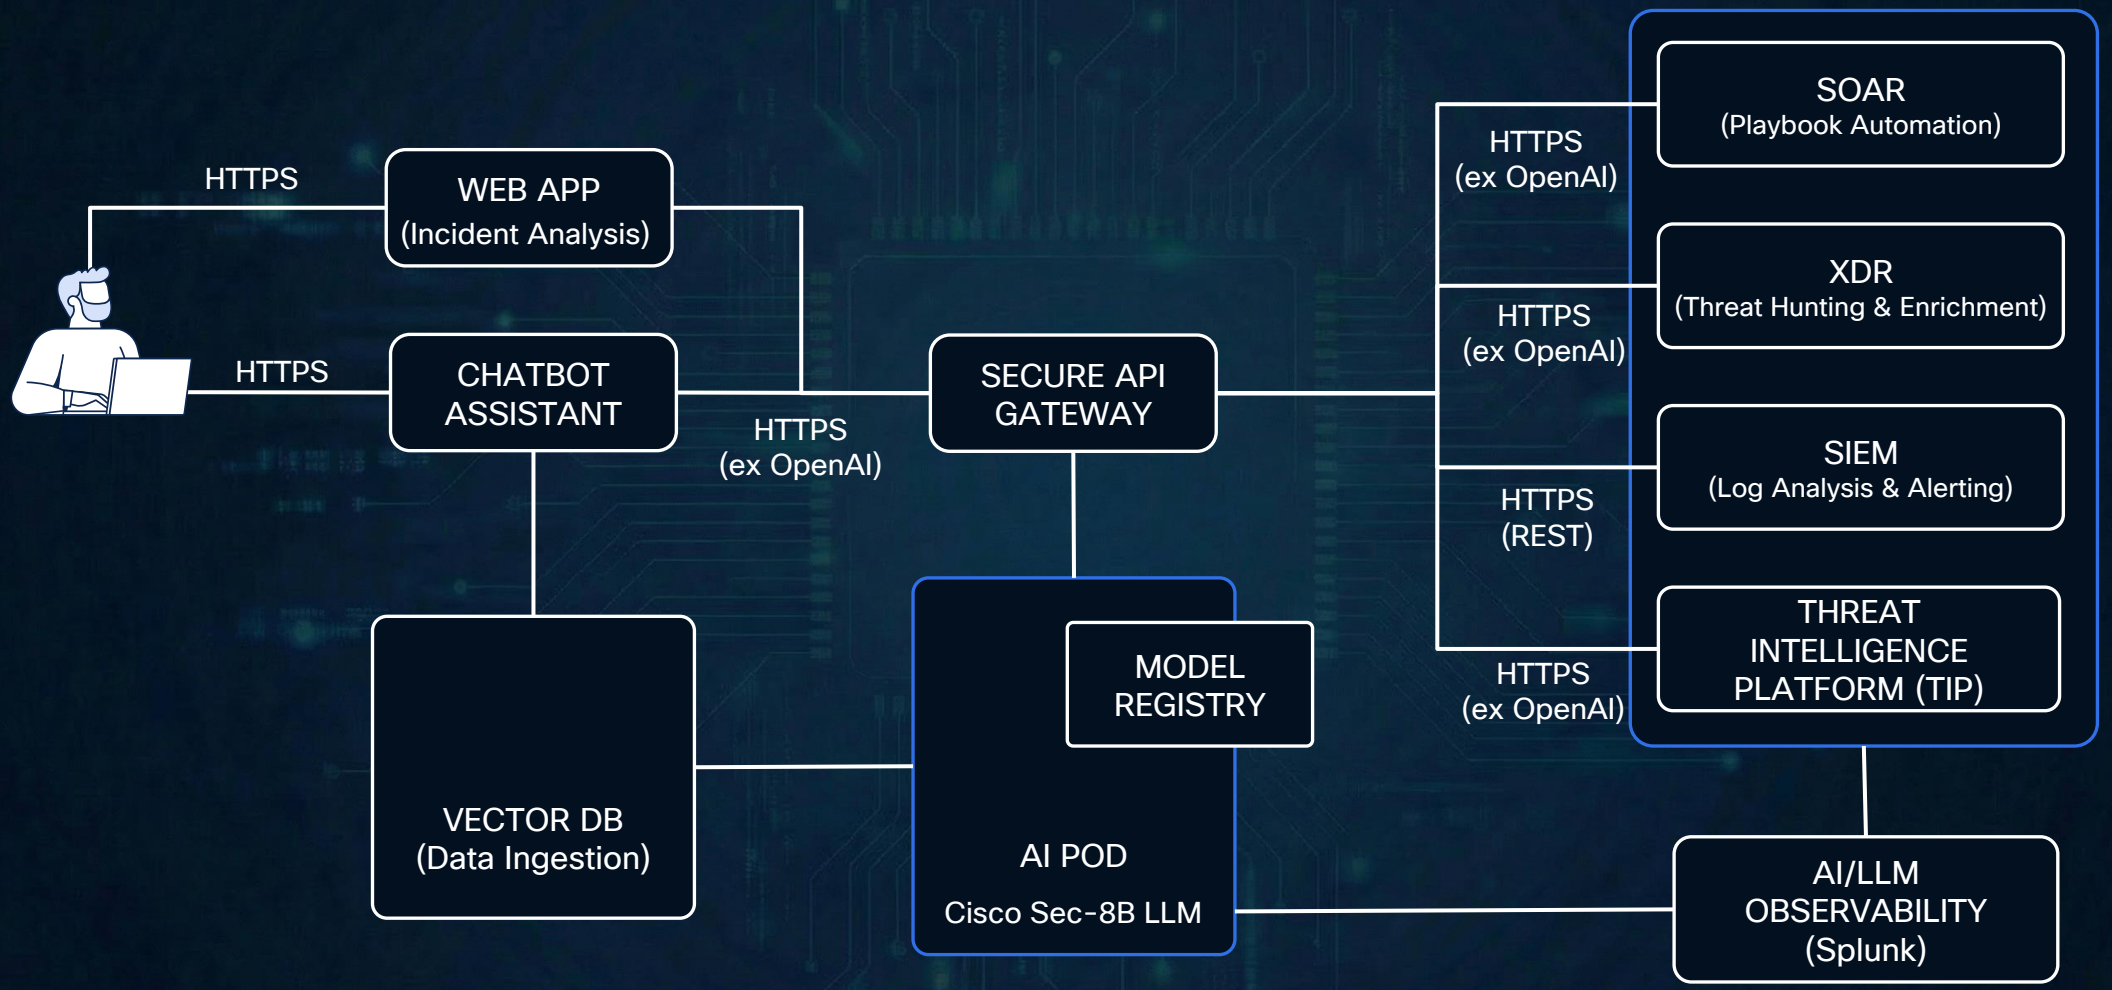

LLM for SOC: Sample Executive Architecture

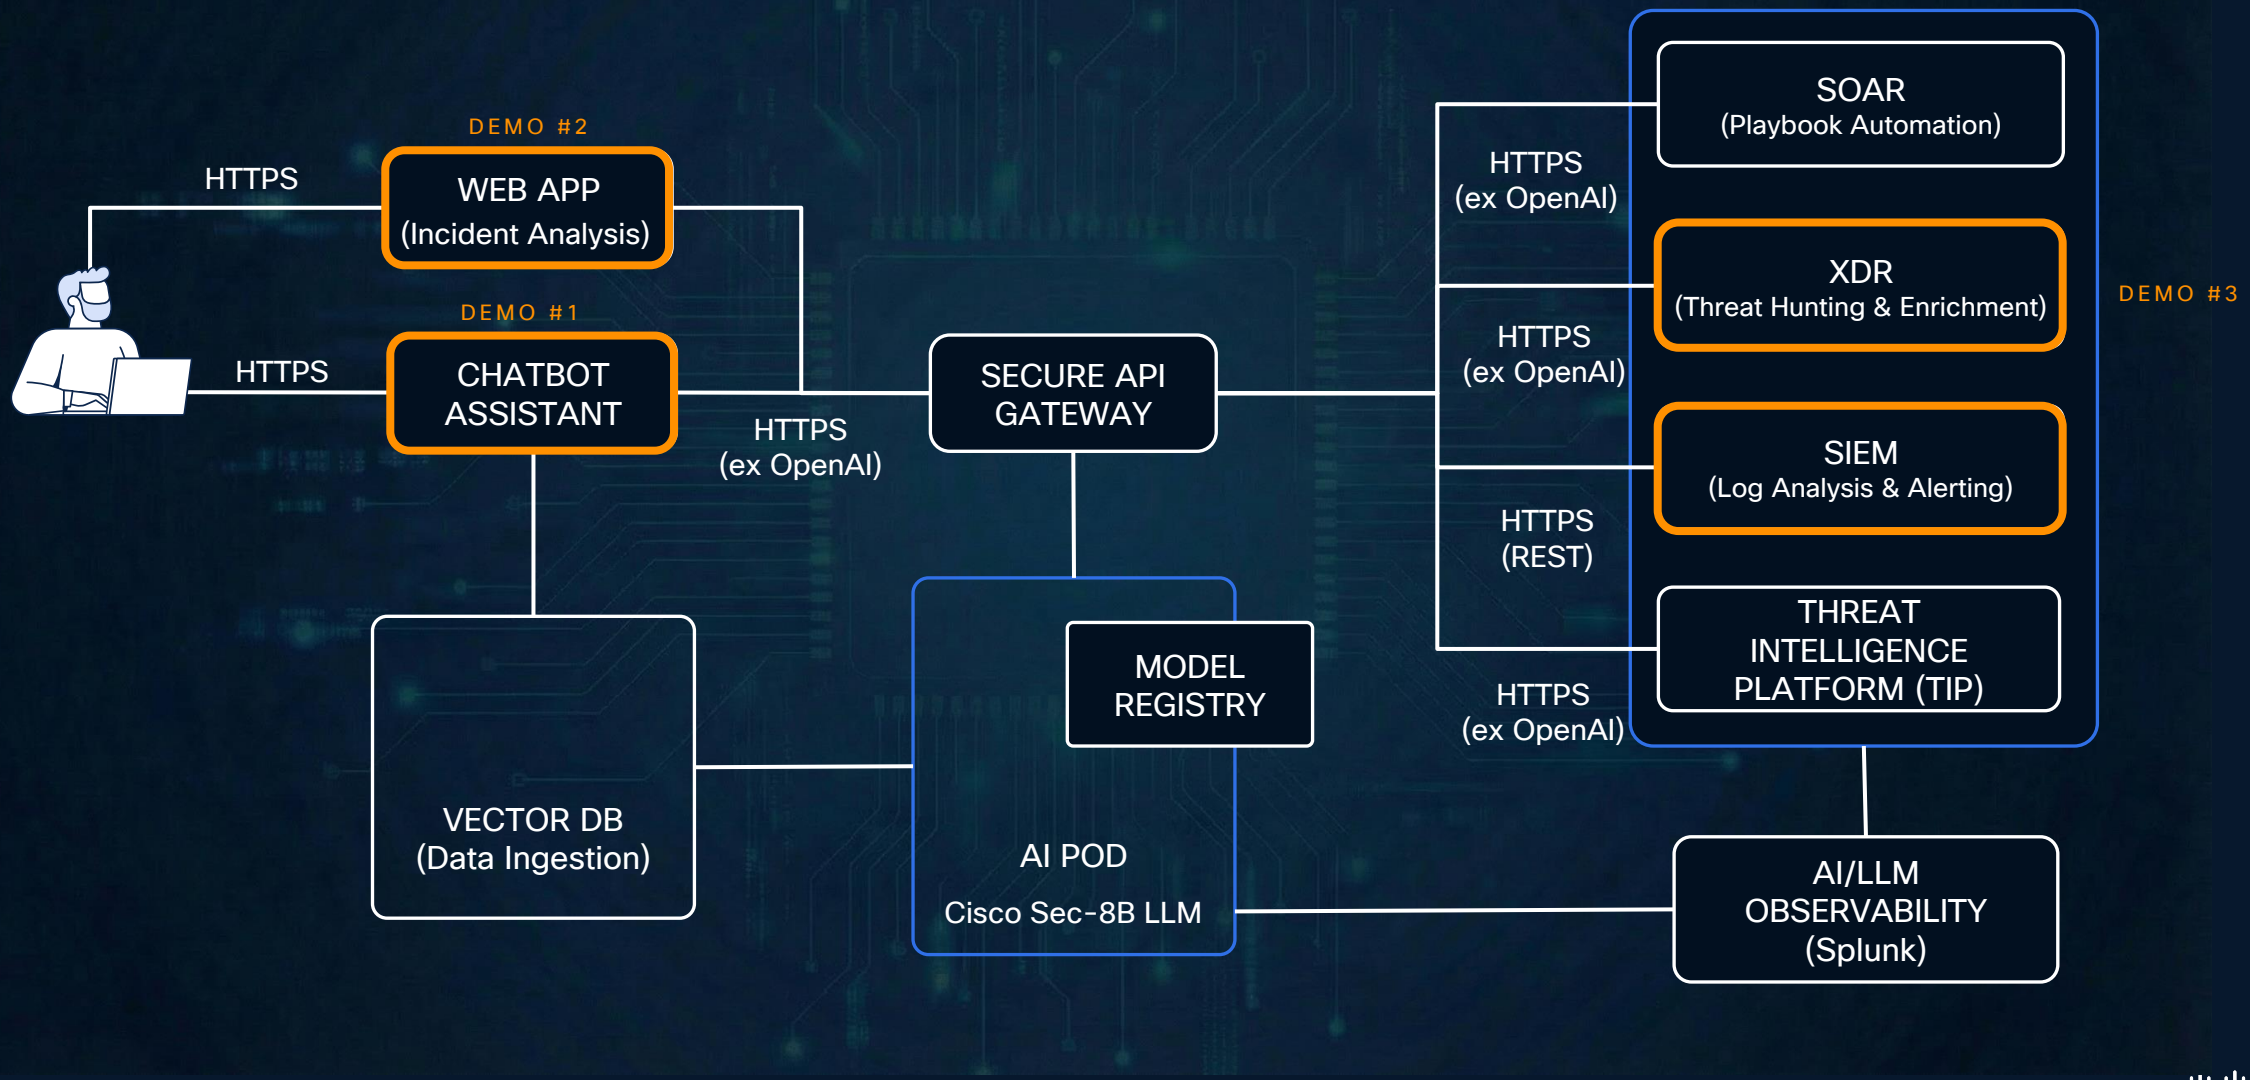

Use Case #1 Demo
Simple Chatbot

Demo Use Case 1

Simple Chatbot

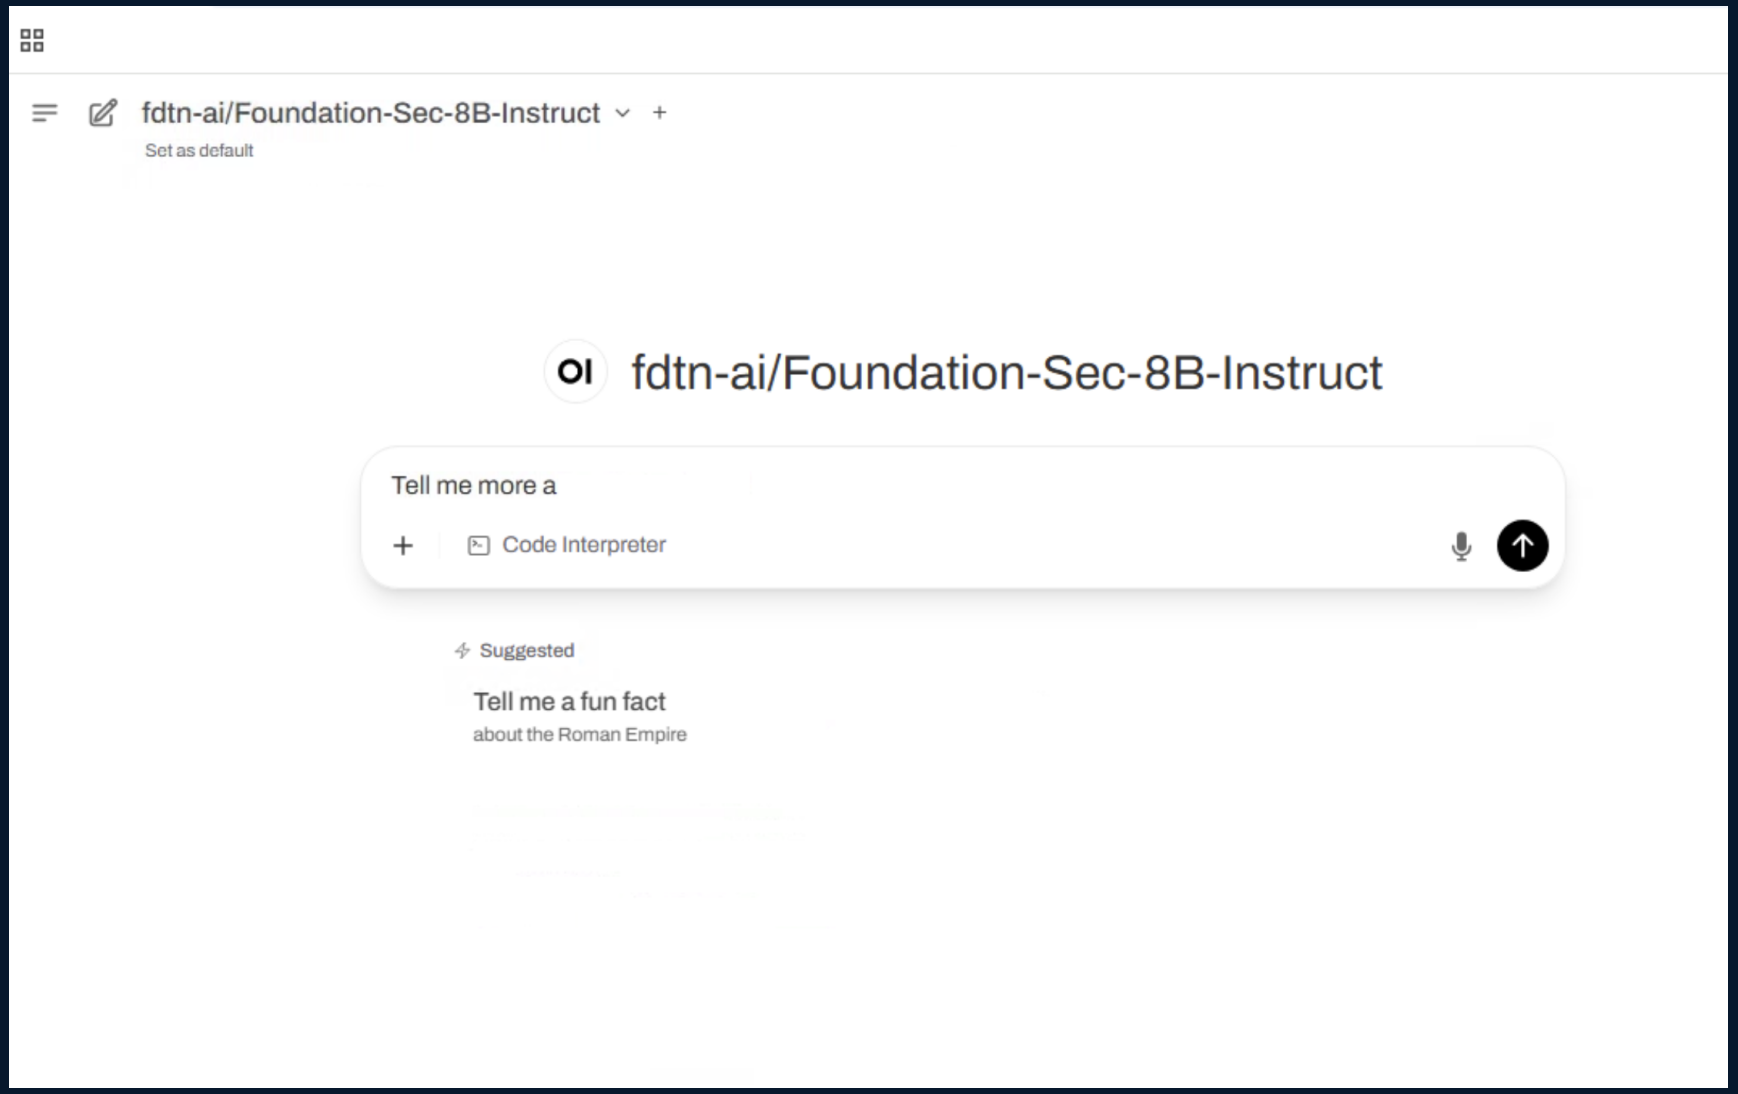

More than just a Chatbot: A Security-Native Interface. Unlike general-purpose LLMs, Sec-8b is pre-loaded with the 'DNA' of cybersecurity.

Use Case #2 Demo
Incident Summarization

Demo Use Case 2

Incident Summarization

Critical cybersecurity incidents often result in lengthy, complex documentation that is difficult to digest, slowing triage, handoffs, and leadership visibility. The goal is to automatically condense these incident reports into concise summaries tailored to analysts or executives – enabling faster response, clearer communication, and more effective decision-making

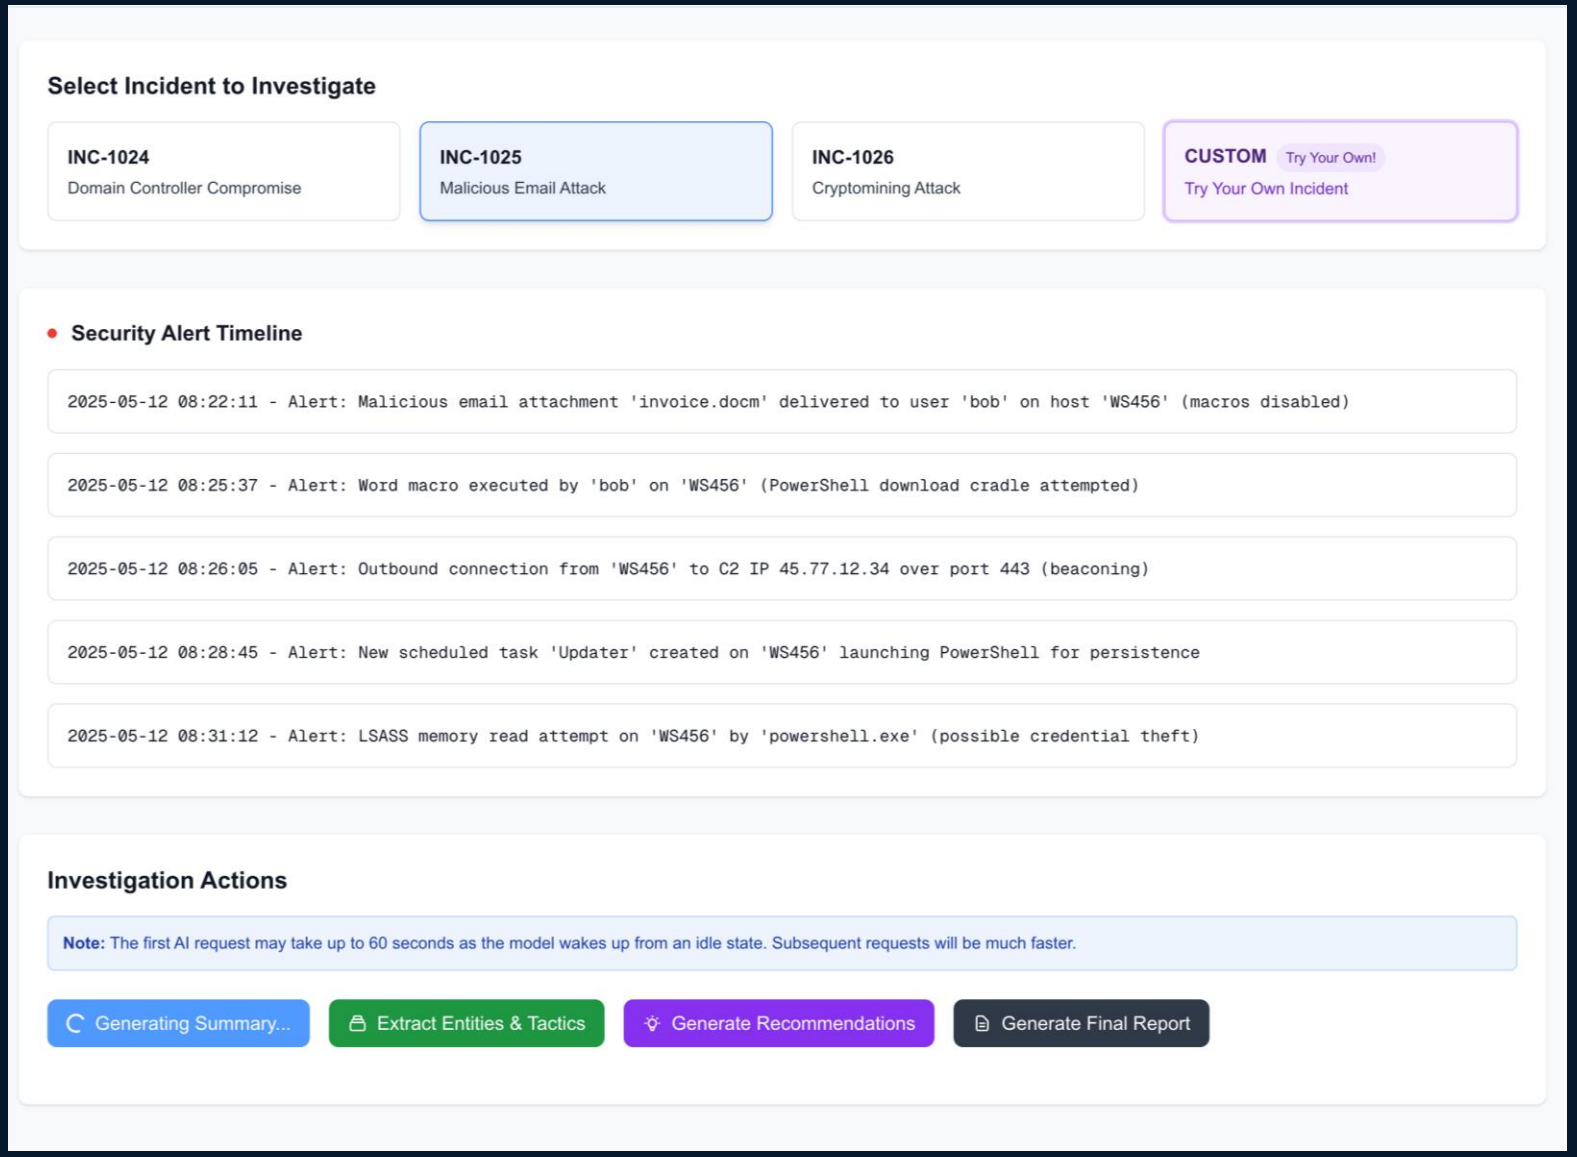

Incident Investigation Dashboard

Blackhat Demo - Cybersecurity Incident Analysis

Select Incident to Investigate

INC-1024
Domain Controller Compromise

INC-1025
Malicious Email Attack

INC-1026
Cryptomining Attack

CUSTOM Try Your Own!
Try Your Own Incident

Security Alert Timeline

```
2025-04-09 10:00:23 - Alert: 5 failed login attempts for user 'alice' on host 'WS123'
```

```
2025-04-09 10:05:10 - Alert: Suspicious PowerShell execution on 'WS123' by 'alice' (malicious script blocked)
```

```
2025-04-09 10:10:45 - Alert: Process dumping LSASS memory on 'WS123' (possible credential theft)
```

```
2025-04-09 10:15:00 - Alert: Successful login of 'alice' to server 'DC1' from host 'WS123'
```

Investigation Actions

Note: The first AI request may take up to 60 seconds as the model wakes up from an idle state. Subsequent requests will be much faster.

Incident Investigation Dashboard

Blackhat Demo - Cybersecurity Incident Analysis

Select Incident to Investigate

INC-1024
Domain Controller Compromise

INC-1025
Malicious Email Attack

INC-1026
Cryptomining Attack

CUSTOM Try Your Own!
Try Your Own Incident

Security Alert Timeline

```
2025-04-09 10:00:23 - Alert: 5 failed login attempts for user 'alice' on host 'WS123'
```

```
2025-04-09 10:05:10 - Alert: Suspicious PowerShell execution on 'WS123' by 'alice' (malicious script blocked)
```

```
2025-04-09 10:10:45 - Alert: Process dumping LSASS memory on 'WS123' (possible credential theft)
```

```
2025-04-09 10:15:00 - Alert: Successful login of 'alice' to server 'DC1' from host 'WS123'
```

Investigation Actions

Note: The first AI request may take up to 60 seconds as the model wakes up from an idle state. Subsequent requests will be much faster.

Use Case #3 Demo
Cisco XDR Integration Example

Demo Use Case 3

Enhancing Security Operations with Cisco XDR and Sec-8b-Instruct LLM

Workflow Integration: dedicated XDR workflow to facilitate API queries to Sec-8B running on Cisco AI POD Foundation-sec compute server, transmitting incident content for analysis.

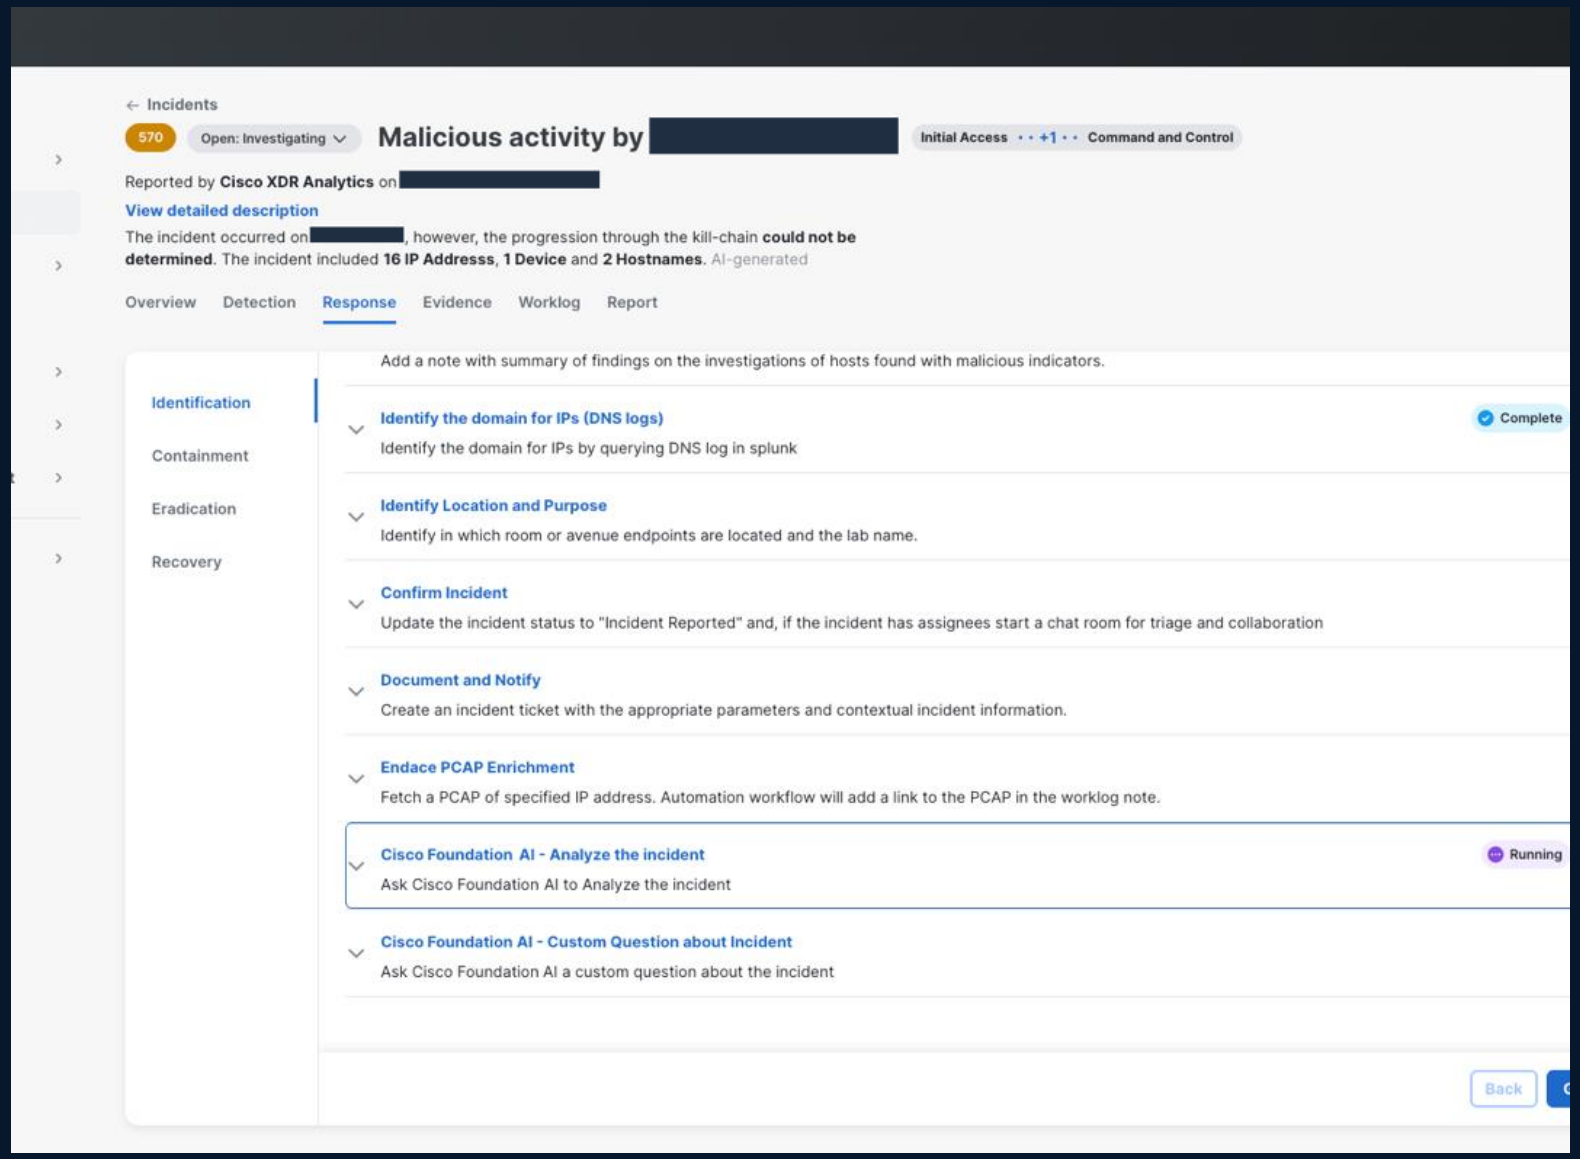

Demo Use Case 3

Workflow Setup

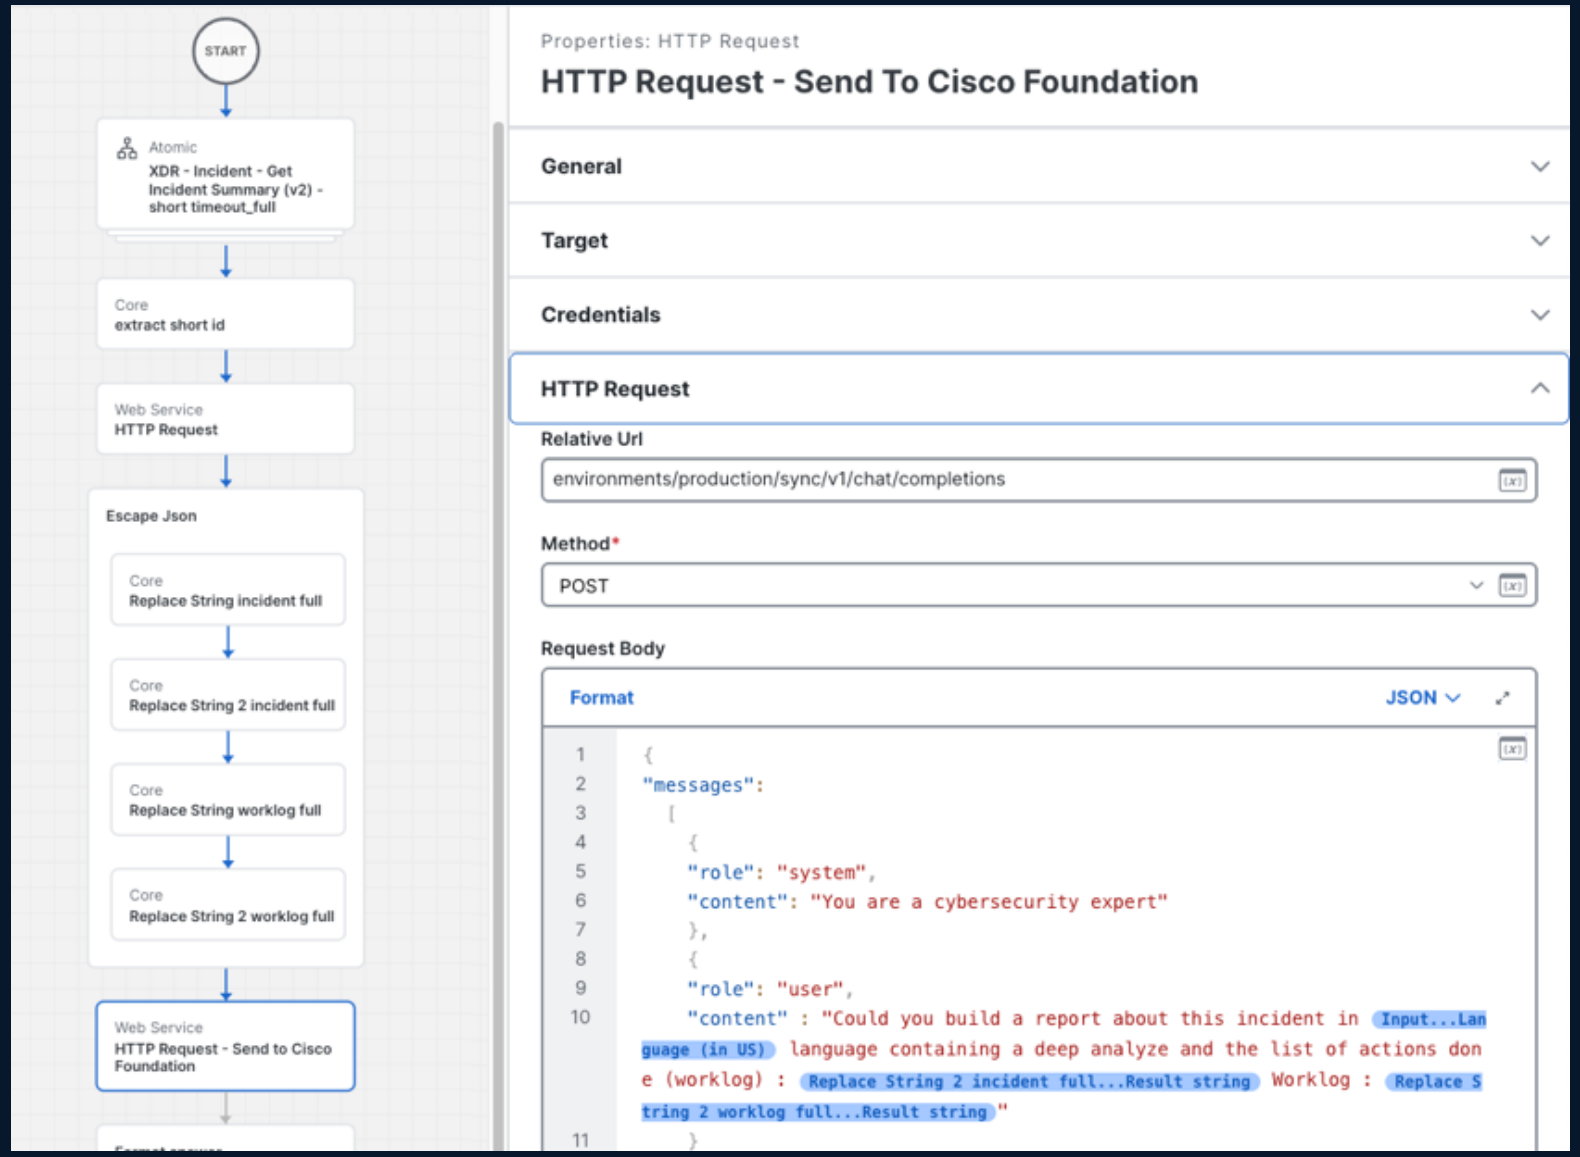

Workflow Integration: A dedicated XDR workflow was established to facilitate API queries to our Foundation-sec model running on Cisco AI POD, transmitting incident content for analysis

Playbook Integration

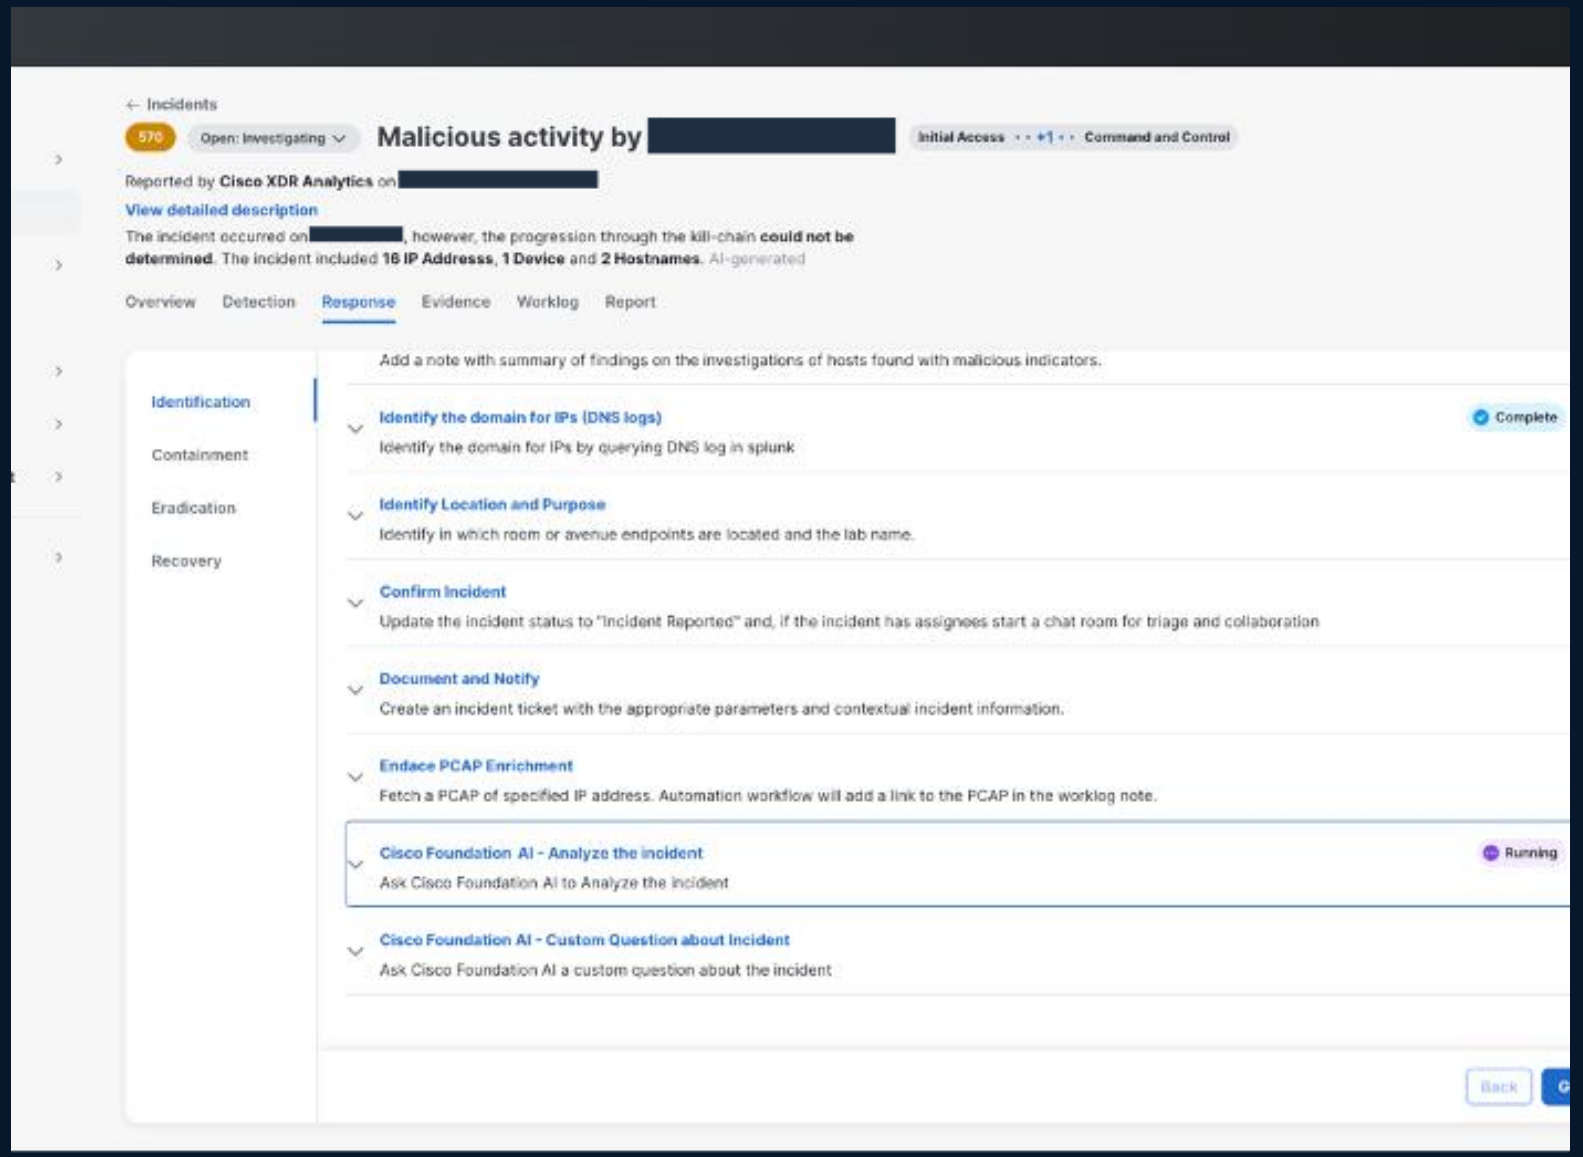

Playbook Integration: The model was further integrated into XDR as an identification playbook. This allowed security analysts to initiate an immediate analysis of any incident by selecting "Ask Cisco Foundation AI to Analyze the incident" directly from the incident view

Control Center
Incidents
Investigate
Intelligence
Automate
Assets
Client Management
Administration

Incidents

197 Incidents | **11** New Incidents | **1** Open Incident | **15** Unassigned Incidents

Search | Last year | Hide closed incidents | Filters | 5 results | Reset all

1 Applied Filter

Priority	Name	Sources	Created	Assigned	Status
210	Malicious IP activity by 10[.]220[.]250[.]37	Cisco Secure Firewall via Splunk	2025-12-10T09:33:22.010Z	IB JB	Closed: Confirmed Threat
210	Malicious activity associated with 195[.]58[.]49[.]249	Cisco Secure Firewall via Splunk	2025-12-10T08:37:11.303Z	JB	Closed: Confirmed Threat
210	Cisco SFW - IP Malicious	Cisco Secure Firewall via Splunk	2025-12-10T08:27:14.803Z	JB	Closed: Confirmed Threat
210	Malicious IP activity by 195[.]58[.]49[.]249	Cisco Secure Firewall via Splunk	2025-12-10T07:37:22.694Z	JB	Closed: Confirmed Threat
680	Secure client overflow attempt by SERVER-MYSQL	Cisco Secure Firewall via Splunk	2025-12-08T09:57:41.769Z	JB	Closed: Confirmed Threat

Use Case #4 Demo
Zero-Shot Security Classification

Demo Use Case 4

Zero-Shot Security Classification

Allows security teams to classify and respond to threats more rapidly and effectively by using advanced AI models that generalize to new, unseen security events without needing extensive retraining or labeled data.

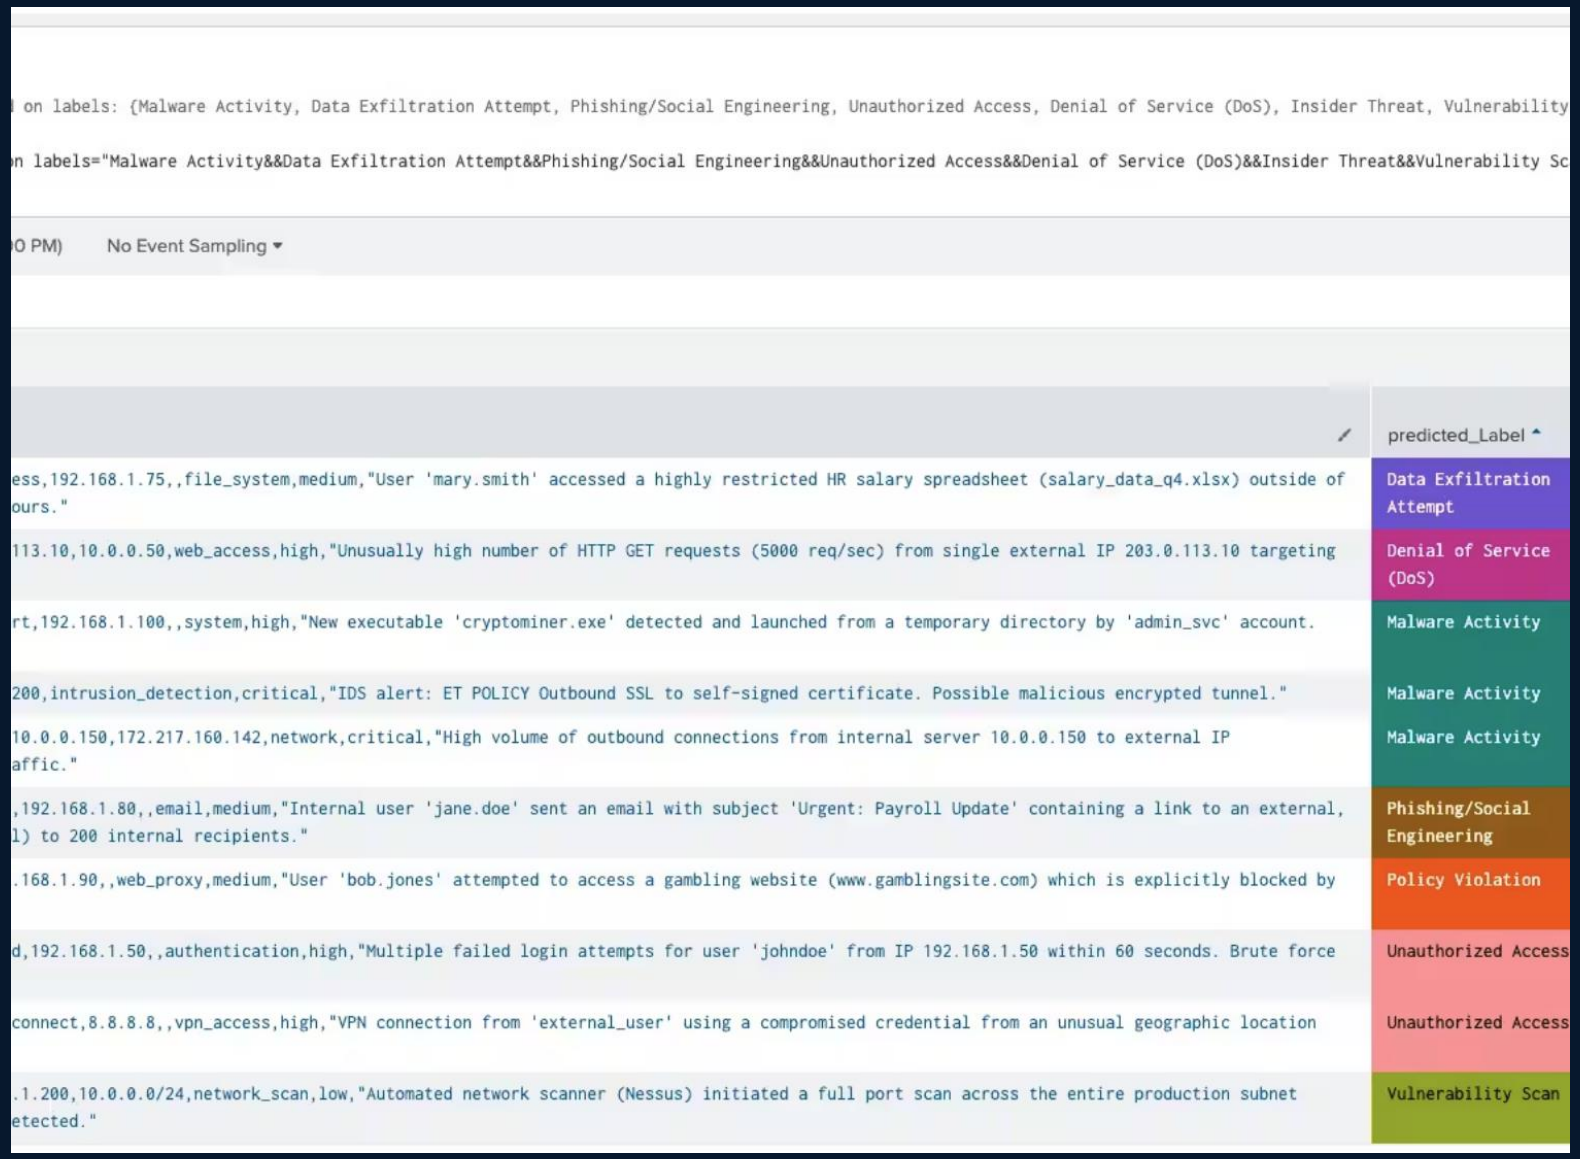

Closing and Q&A

Empower Your SOC with AI-Driven Security

Experience Sec-8b:
- Free to try, easy to run on your laptop, or at scale

Enterprise-Ready Deployment:
- Run on Cisco Secure AI Factory; delivers unmatched reliability and security
- Flexible GPU support—NVIDIA or AMD

Partner with Cisco:
- Experiment, innovate and strengthen your defense
- Cisco experts are ready to support your AI security journey

Stay ahead—secure your future with Cisco AI-first solutions

Where to go next?

This cookbook repository helps you get started with using models developed by Cisco Foundation AI team.
It will cover many security related use cases, deployment options, practical adaption approaches like finetuning and so forth.

Complete your session surveys

Complete your surveys in the Cisco Events App.

Complete a minimum of 4 session surveys and the overall event survey to receive a unique Cisco Live t-shirt.

(from 11:30 on Thursday, while supplies last)

Continue your education

Visit the Cisco Showcase for related demos

Book your one-on-one Meet the Engineer meeting

Visit the Technical Solutions Clinics to discuss your technical questions

Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs

Visit the On-Demand Library for more sessions at CiscoLive.com/On-Demand

Contact us at: damaggi@cisco.com, sgioia@cisco.com

Thank you


CISCO