Enhancing DevOps with MLOps and MLSecOps- Guardrails around AI powered Applications

Jatin Sachdeva
Principal Security Architect

cisco Live !

Cisco Webex App

Questions?

Use Cisco Webex App to chat
with the speaker after the session

How

1. Find this session in the Cisco Live Mobile 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 June 13, 2025.

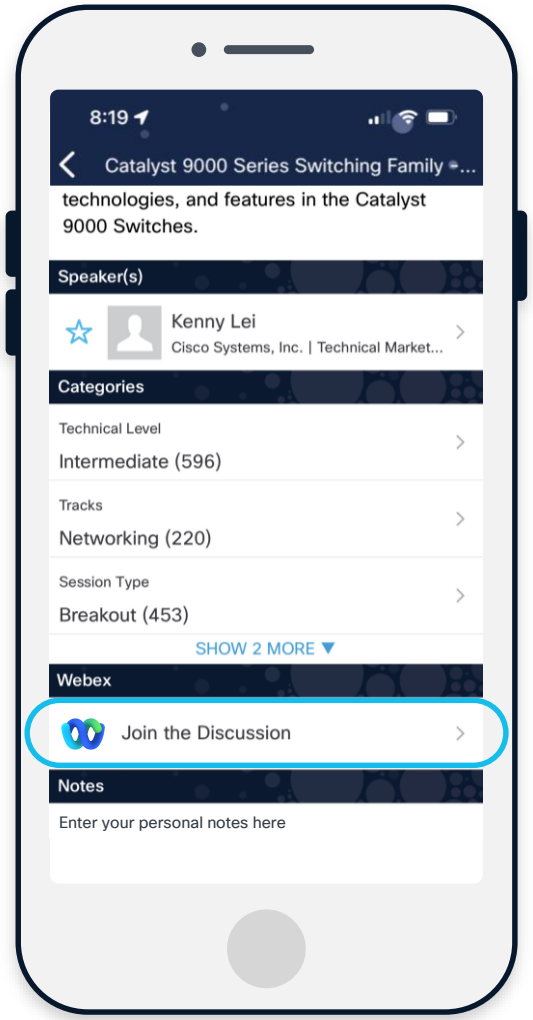

https://ciscolive.ciscoevents.com/
ciscolivebot/**#BRKCLD-1006**

Agenda

01 Intro - What is MLOps?

02 A sample LLM application

03 Where are the Threats!

04 Securing AI with Cisco

05 Outro

About your speaker

Fun fact – I am Indian, but I am not into cricket or spicy food!

Play

- Live in Melbourne, Australia with my lovely family and this crazy fellow
- Can nerd out on anything from tech and cars to fitness and nutrition!

Work

- 23 years in security industry, 20 in Cisco.
- Knowledge seeker hence certs a plenty – CISSP, CISA, CEH, GWAPT, GSEC, GCIA, GCIH, GCSA, GPCS, SFCE, Associate C|CISO
- Prior to Cisco – security consulting, implementation and audit

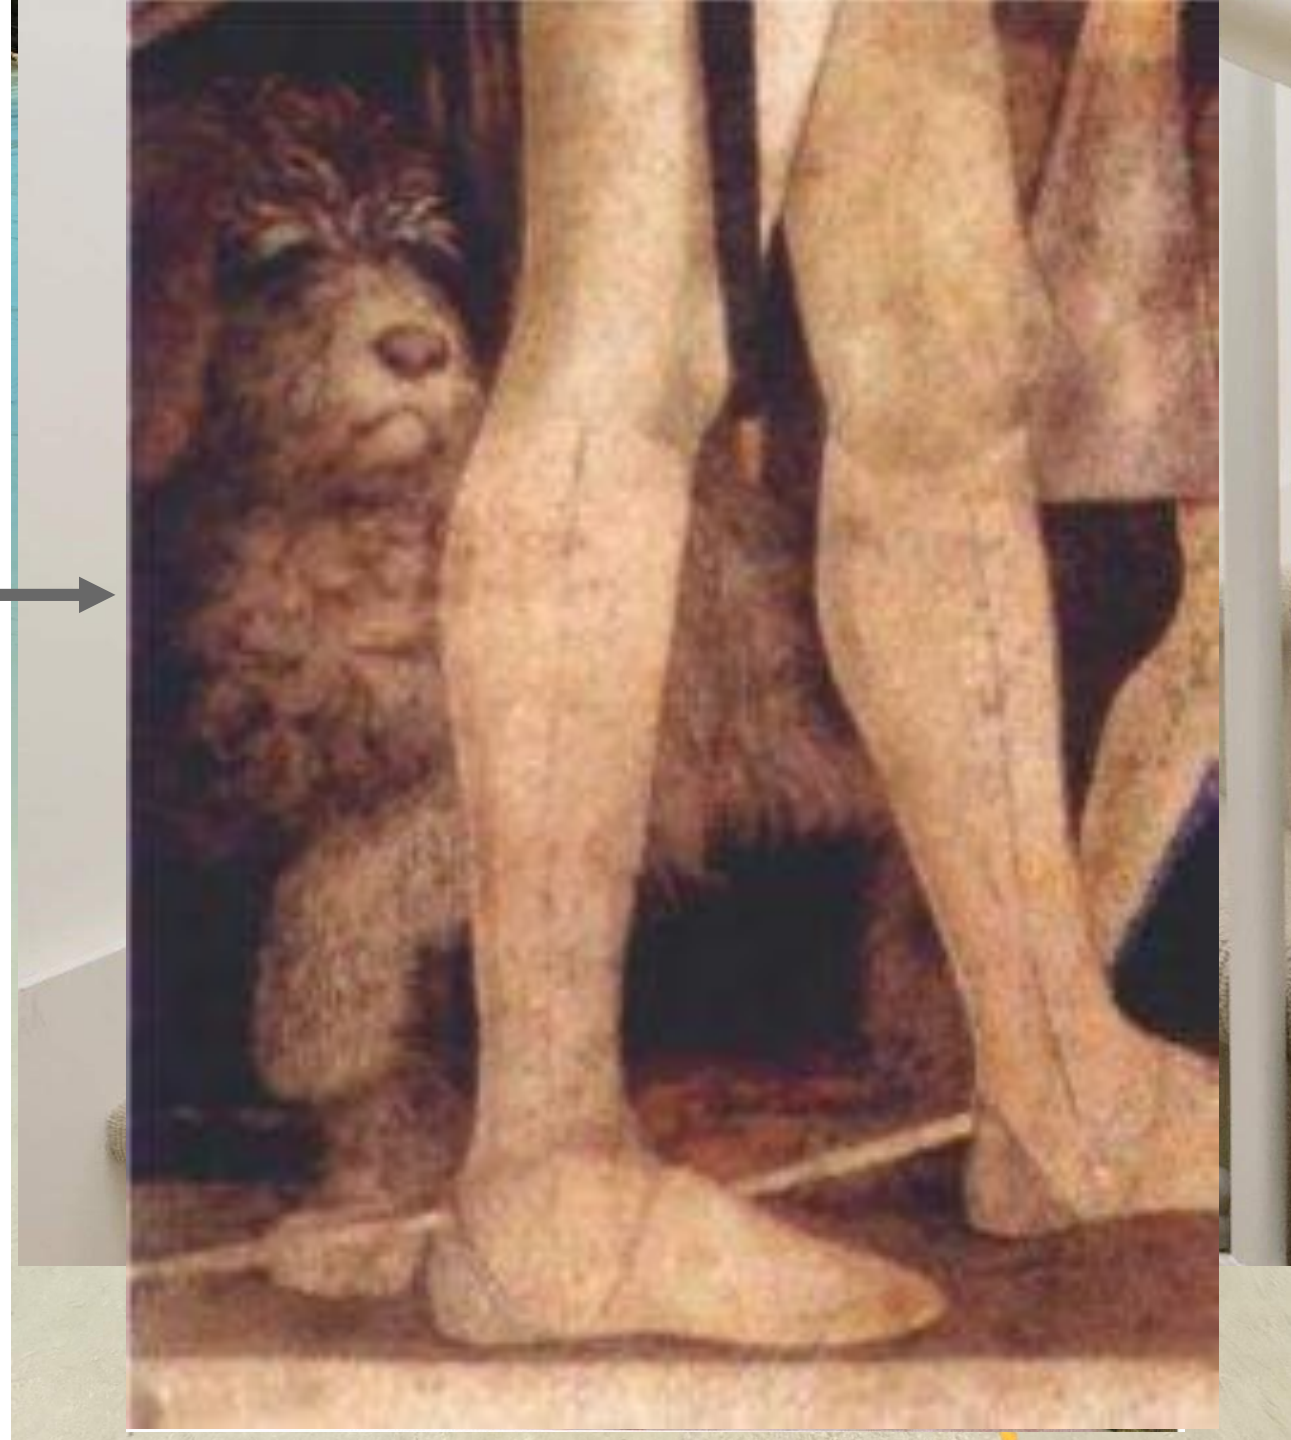

Let’s start with a customer facing application

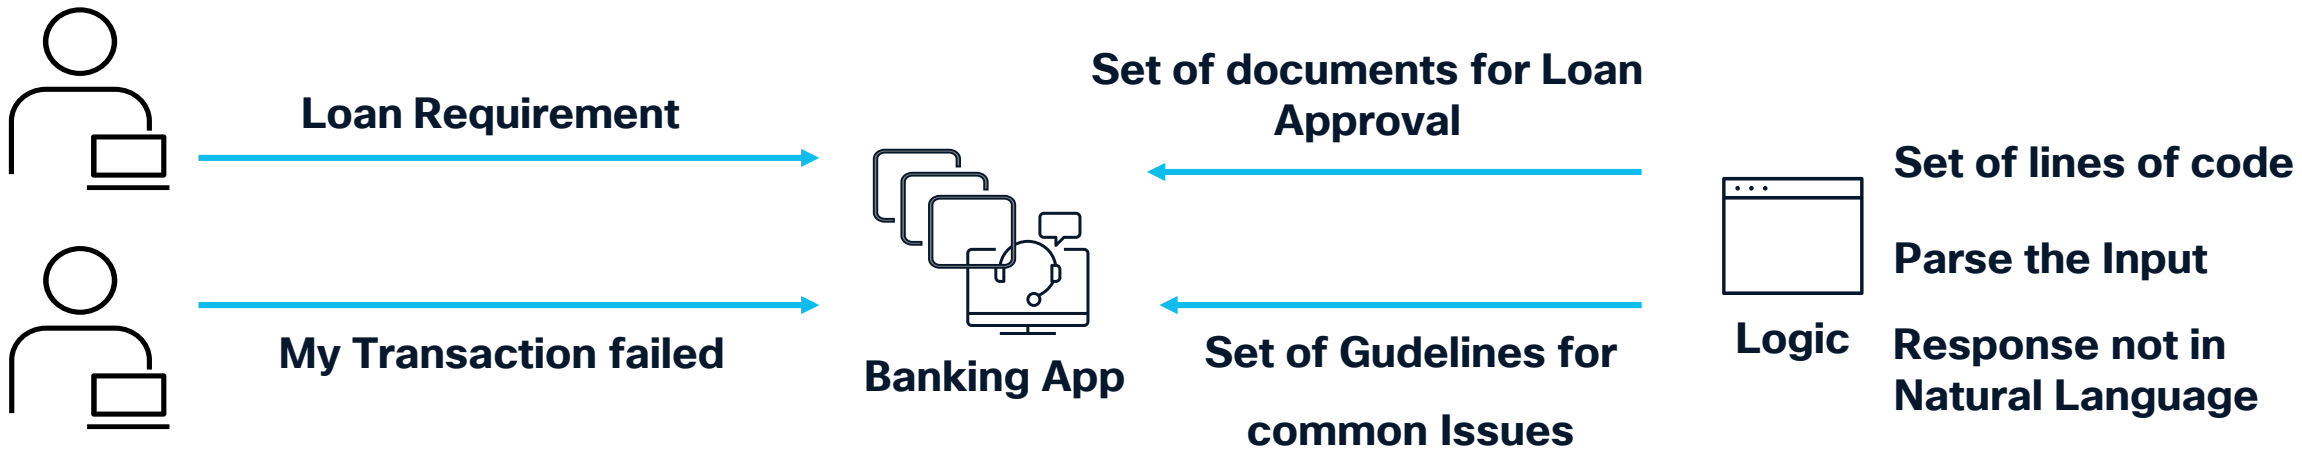

Now let’s add AI Assistance to it

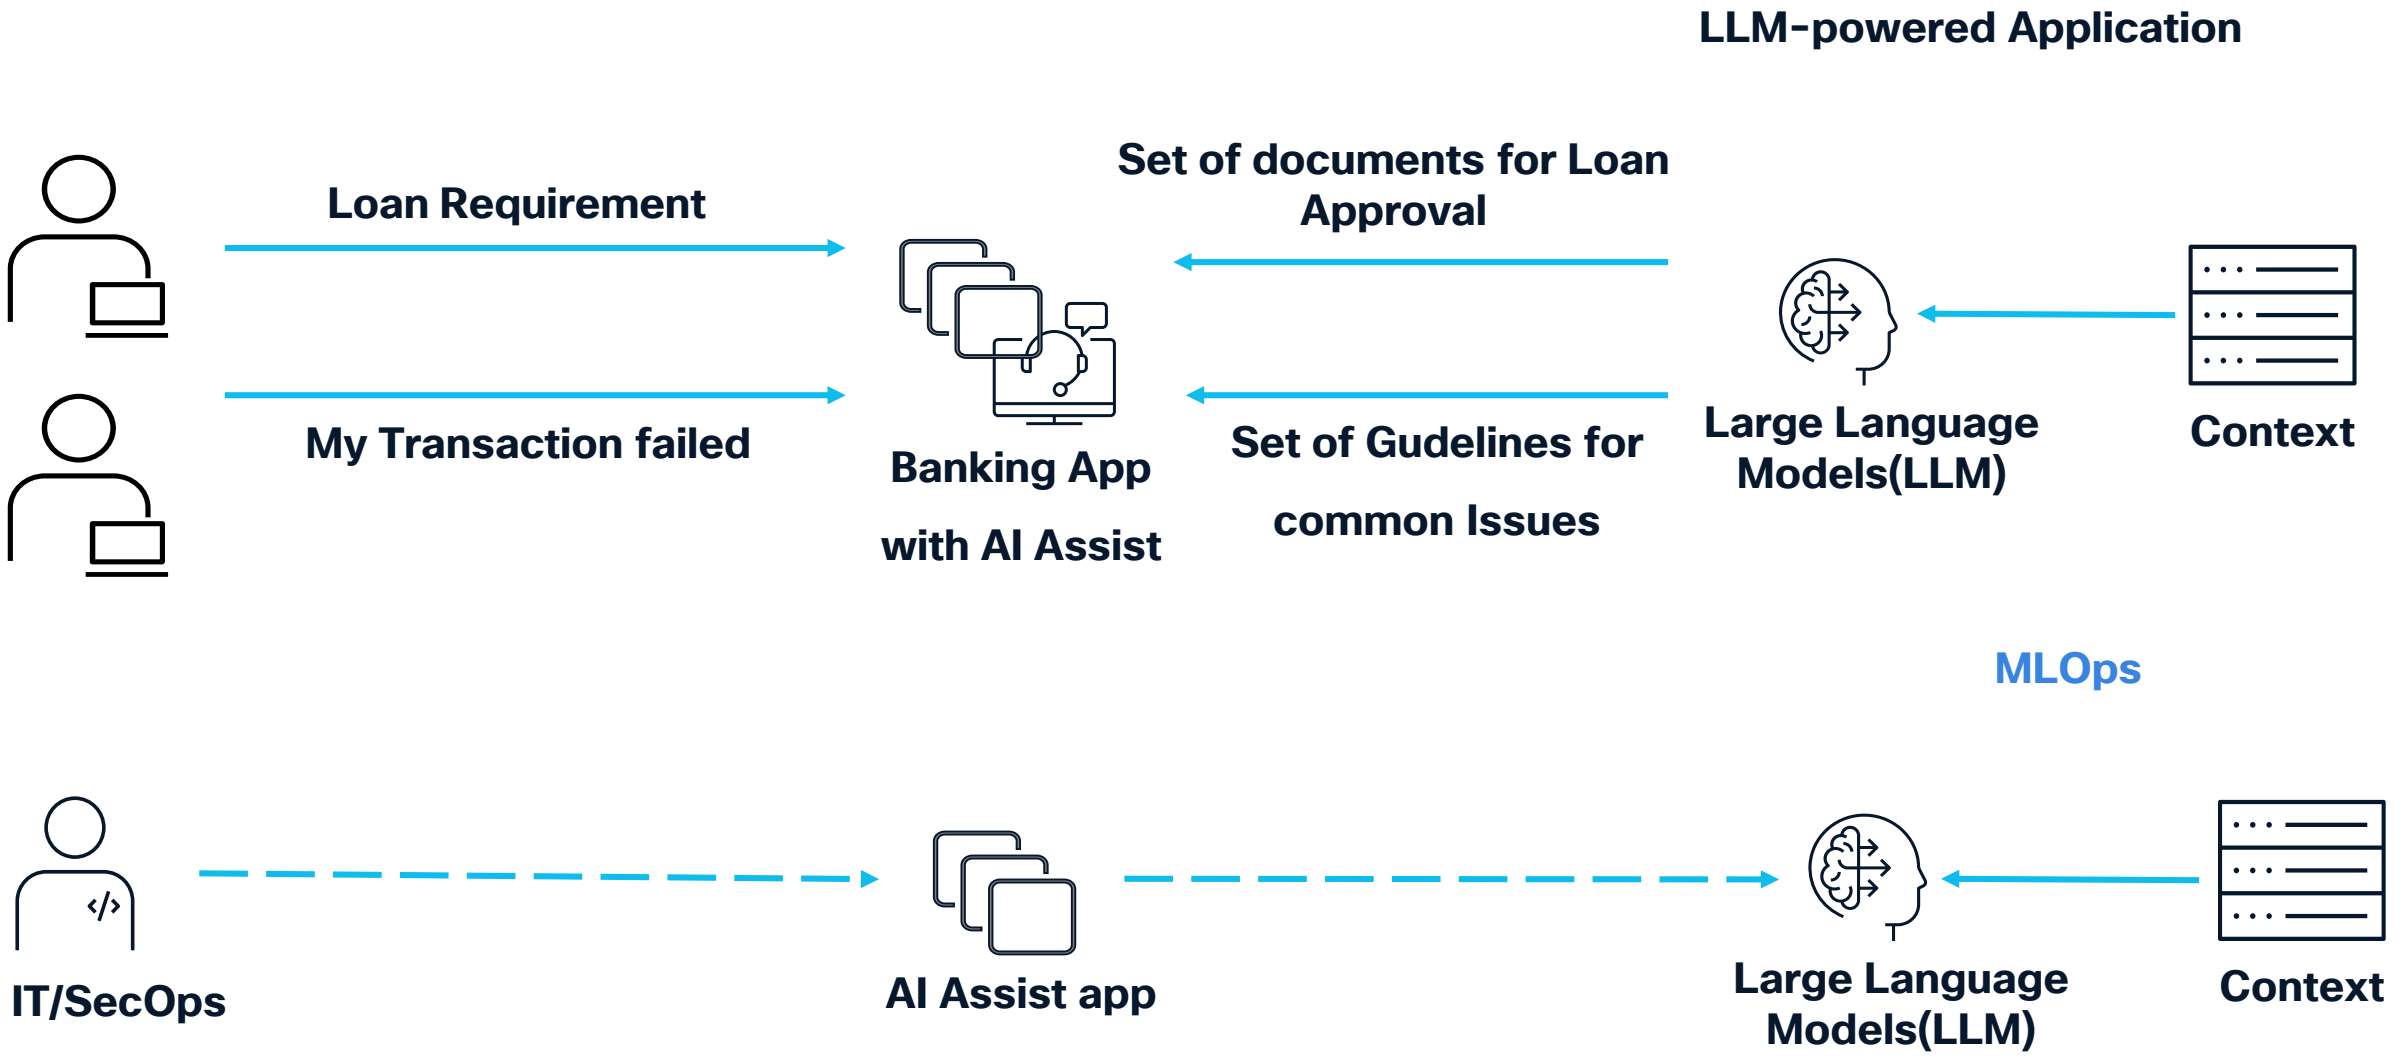

DevOps -> MLOps

Application Lifecycle

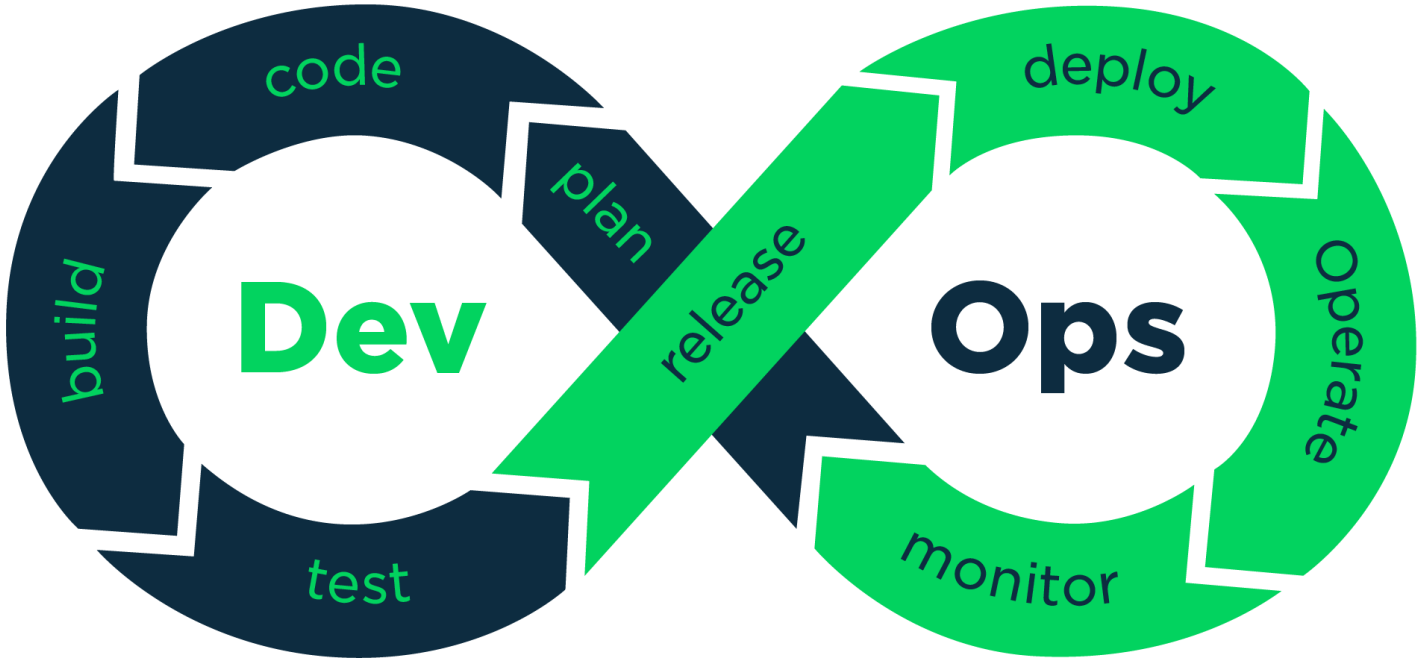

Adding LLM requirements into Application Lifecycle

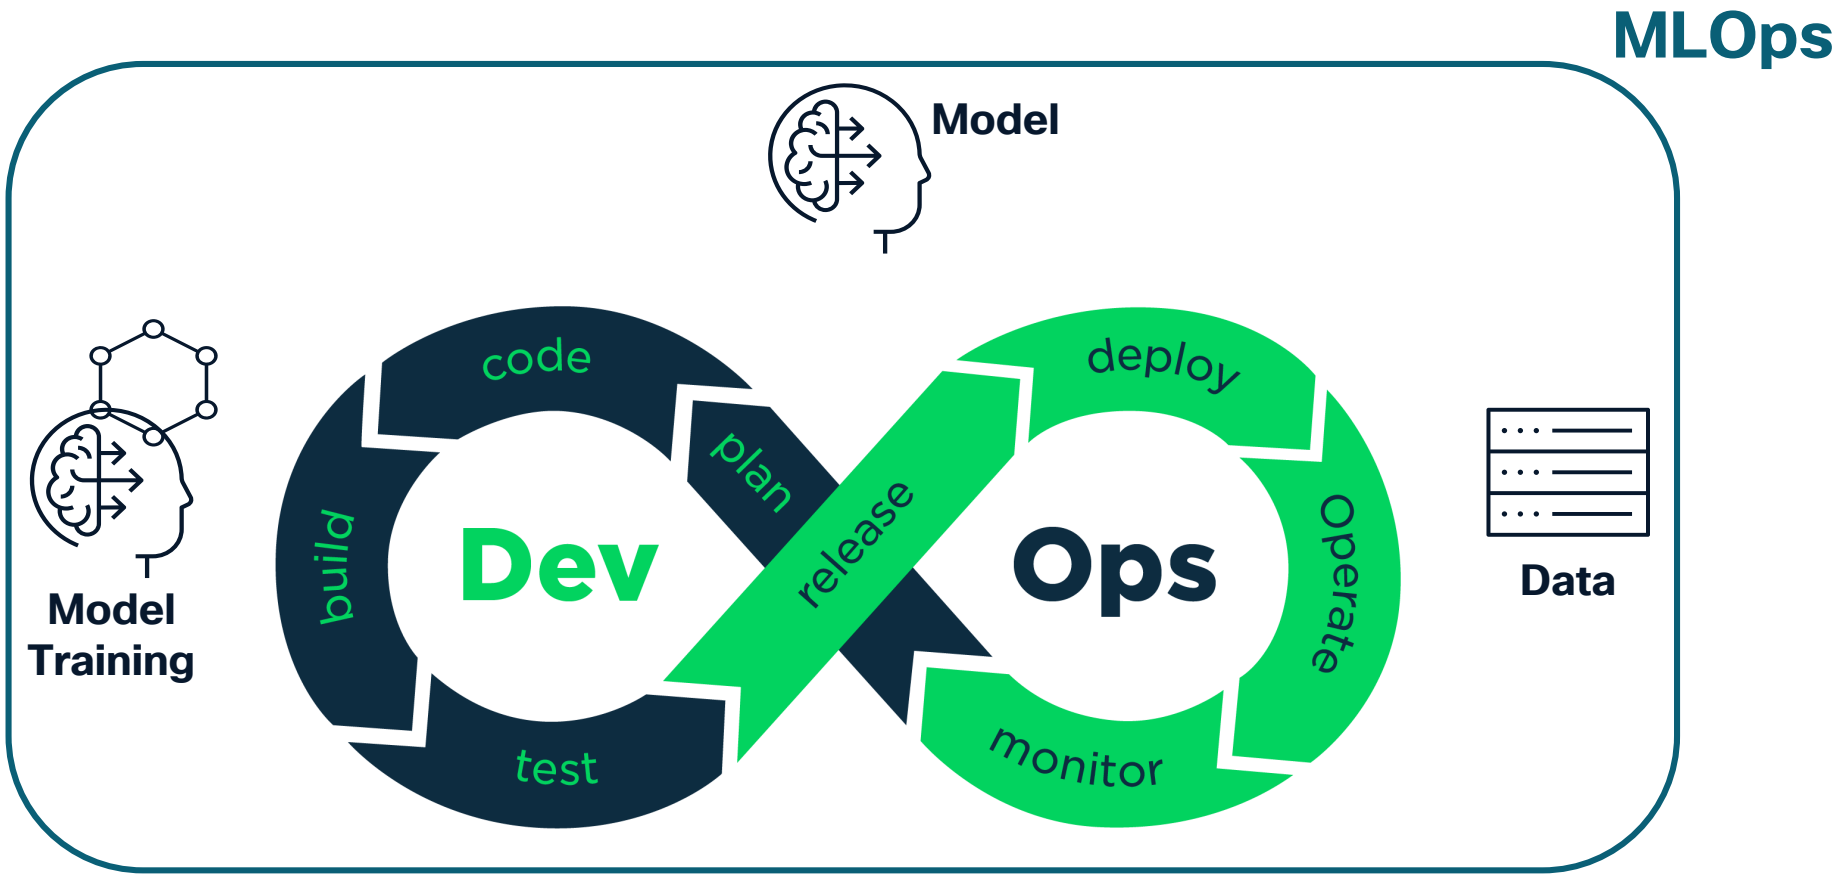

MLOps = DataOps + DevOps + ModelOps

MLOps Core Components

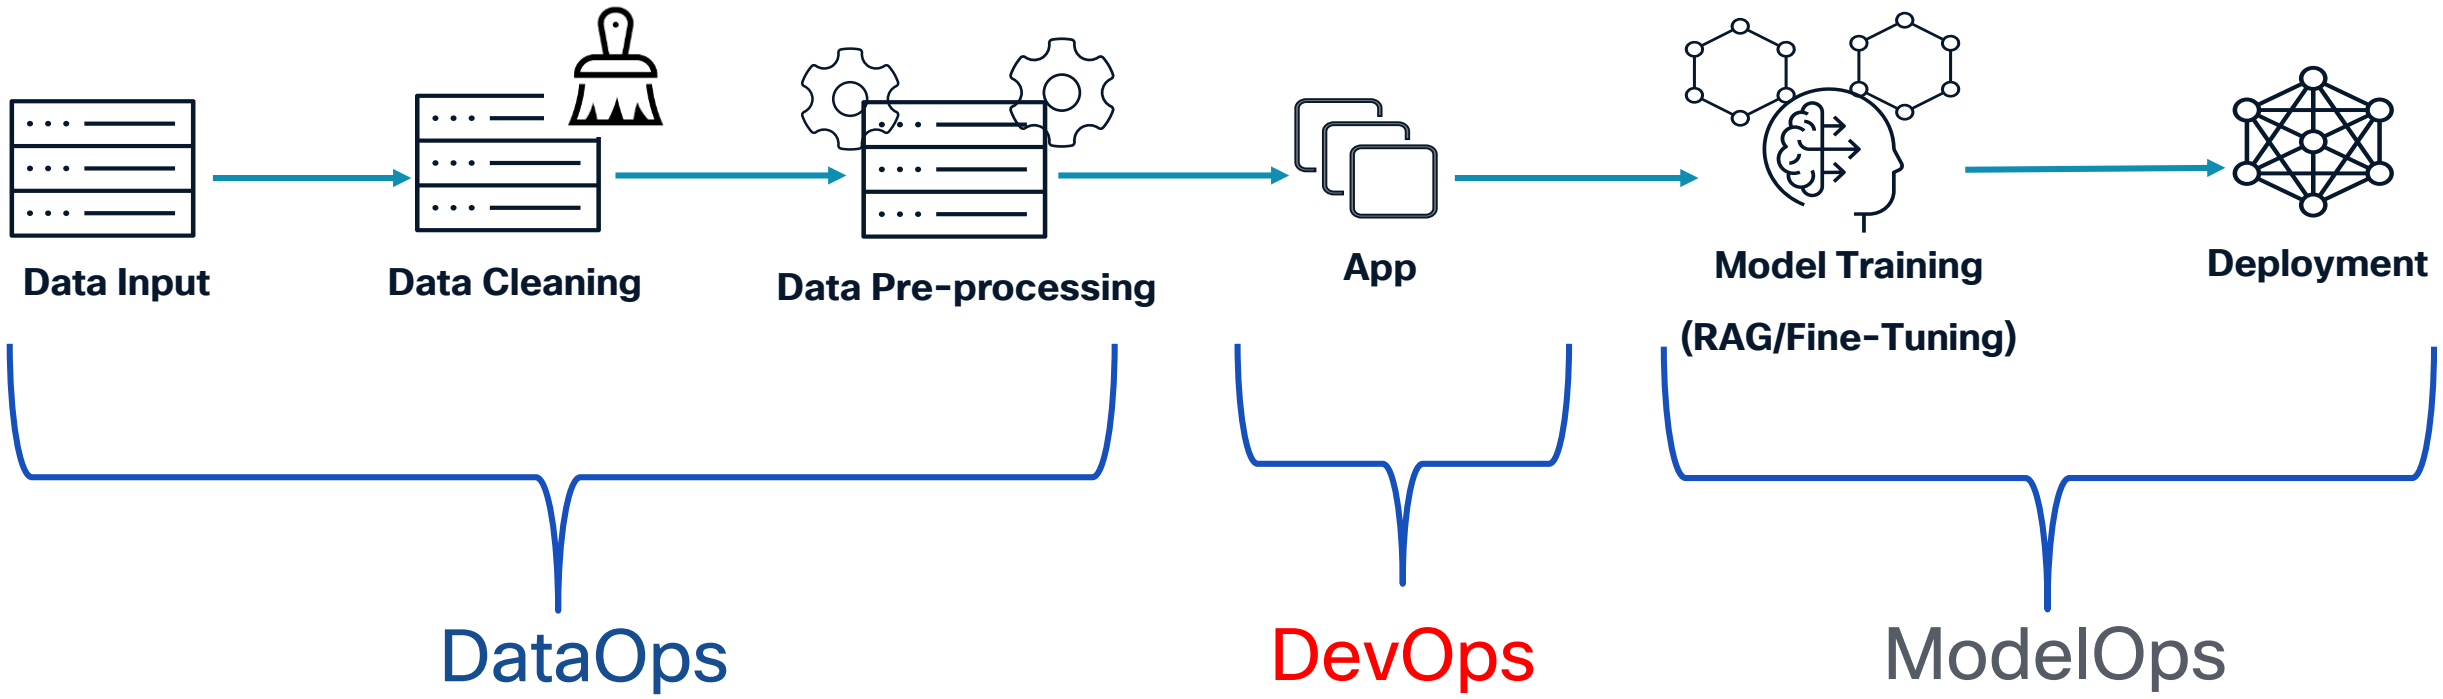

DataOps

Data Input

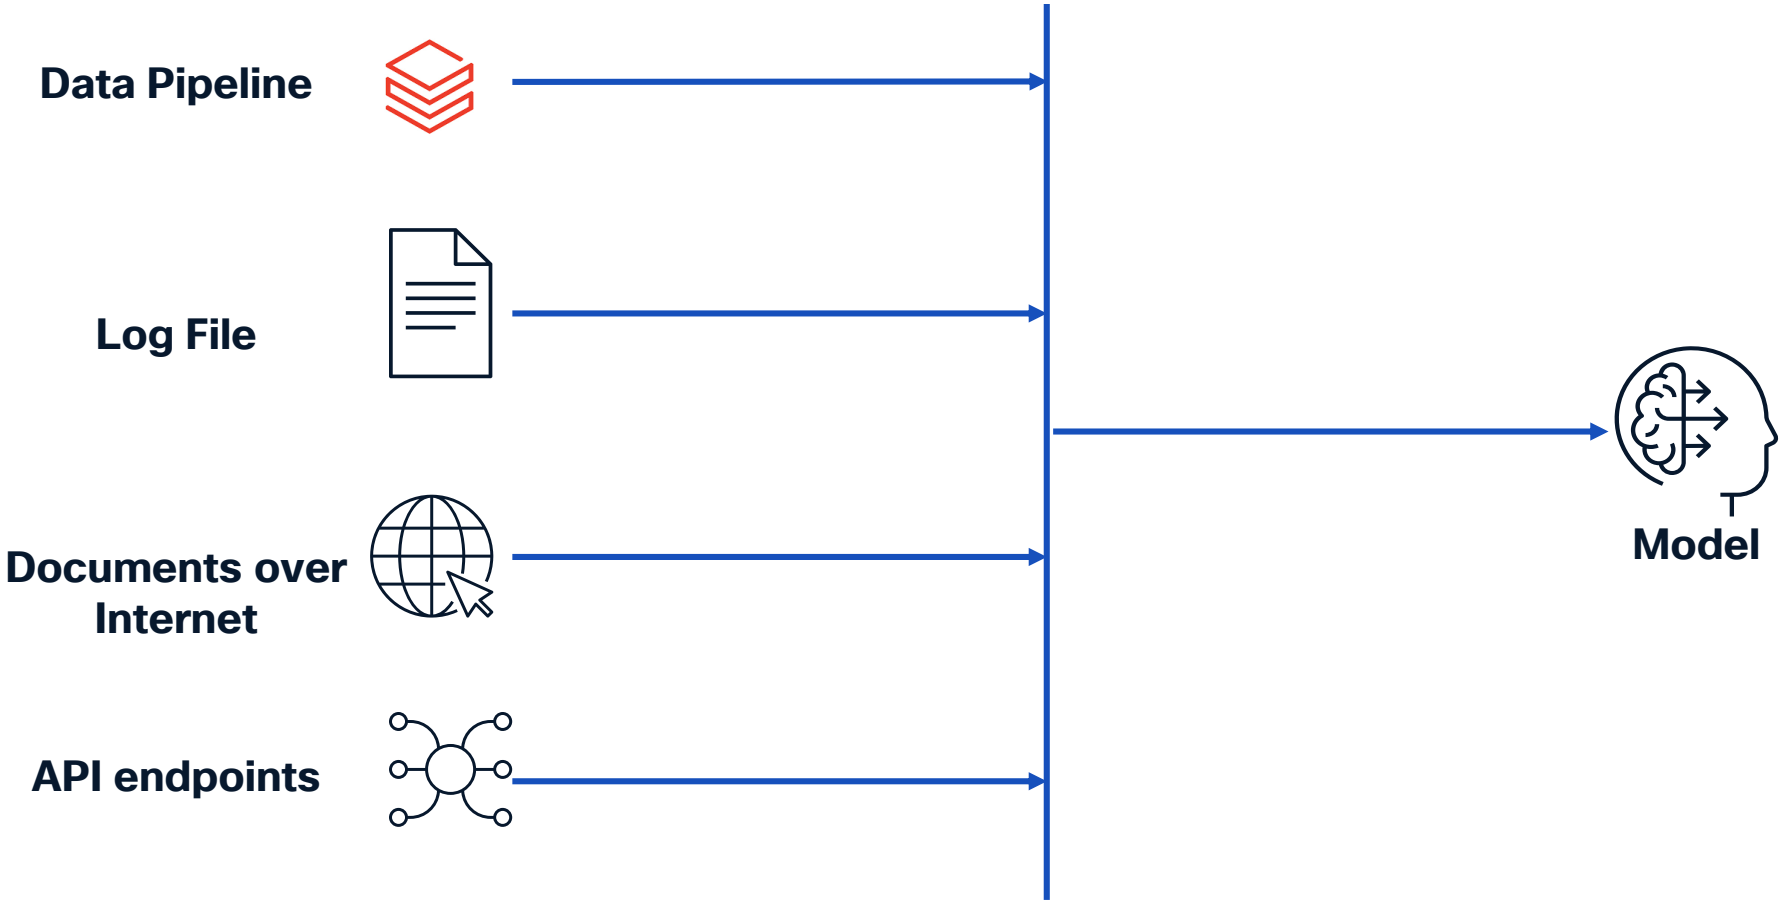

Data Cleaning Example – Handling log files

2025-04-01 09:59:12 [5568] [WARN] < 14> Dns Protection **IPv4** State Machine: Error: Can only have one disconnected session - we have 2 - not using user specific policies

2025-04-01 09:59:12 [5568] [WARN] < 15> Dns Protection **IPv6** State Machine: Error: Can only have one disconnected session - we have 2 - not using user specific policies

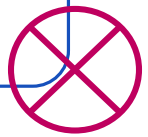

2025-03-28 16:16:12 [8976] [**WARN**] < 10> Device Registration: Exception occurred while deserializing device registration response message; there was likely a problem with the communications channel. Registration will be reattempted later.
System.Runtime.Serialization.SerializationException: There was an error deserializing the object of type Core.RoamingDeviceCreateResponseData. Encountered unexpected character 'T'. ---> System.Xml.XmlException: Encountered unexpected character 'T'.

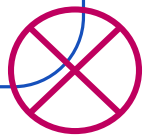

Data Pre-processing example – convert tables to JSON

Management Center Version	Oldest Device Version You Can Manage
7.7	7.2
7.6	7.1
7.4 Last support for NGIPS device management.	7.0
7.3	6.7
7.2	6.6
7.1	6.5
7.0	6.4
6.7	6.3
6.6	6.2.3
6.5	6.2.3
6.4	6.1
6.3	6.1
6.2.3	6.1
6.2.2	6.1
6.2.1	6.1
6.2	6.1
6.1	5.4.0.2/5.4.1.1
6.0.1	5.4.0.2/5.4.1.1
6.0	5.4.0.2/5.4.1.1
5.4.1	5.4.1 for ASA FirePOWER on the ASA-5506-X series, ASA5508-X, and ASA5516-X. 5.3.1 for ASA FirePOWER on the ASA5512-X, ASA5515-X, ASA5525-X, ASA5545-X, ASA5555-X, and ASA-5585-X series. 5.3.0 for Firepower 7000/8000 series and legacy devices.

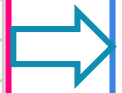

Management Center Version Oldest Device Version You Can Manage\n0 7.7 7.2\n1 7.6 7.1\n2 7.4 Last support for NGIPS device management. 7.0\n3 7.3 6.7\n4 7.2 6.6\n5 7.1 6.5\n6 7.0 6.4\n7 6.7 6.3\n8 6.6 6.2.3\n9 6.5 6.2.3\n10 6.4 6.1\n11 6.3 6.1\n12 6.2.3 6.1\n13 6.2.2 6.1\n14 6.2.1 6.1\n15 6.2 6.1\n16 6.1 5.4.0.2/5.4.1.1\n17 6.0.1 5.4.0.2/5.4.1.1\n18 6.0 5.4.0.2/5.4.1.1\n19 5.4.1 5.4.1 for ASA FirePOWER on the ASA-5506-X seri...

DevOps for LLM inclusion

Interactions with LLMs

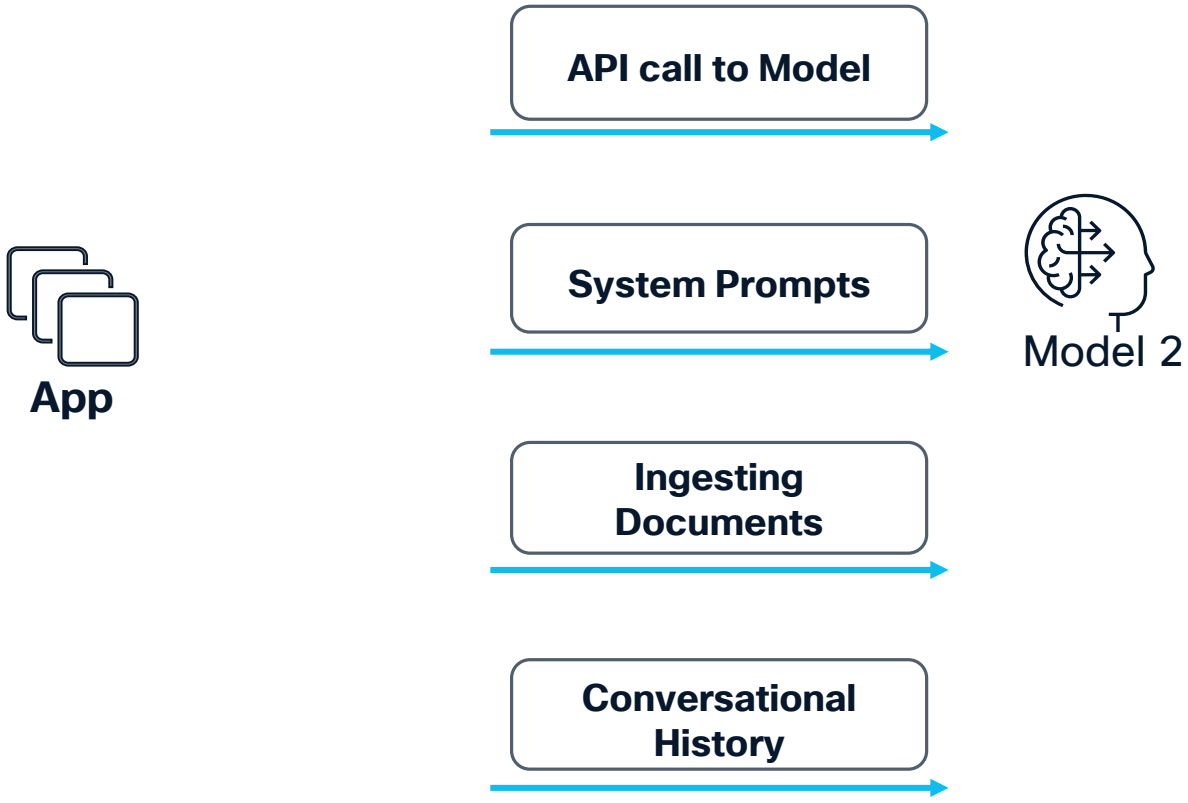

Interactions with LLMs – a modular approach

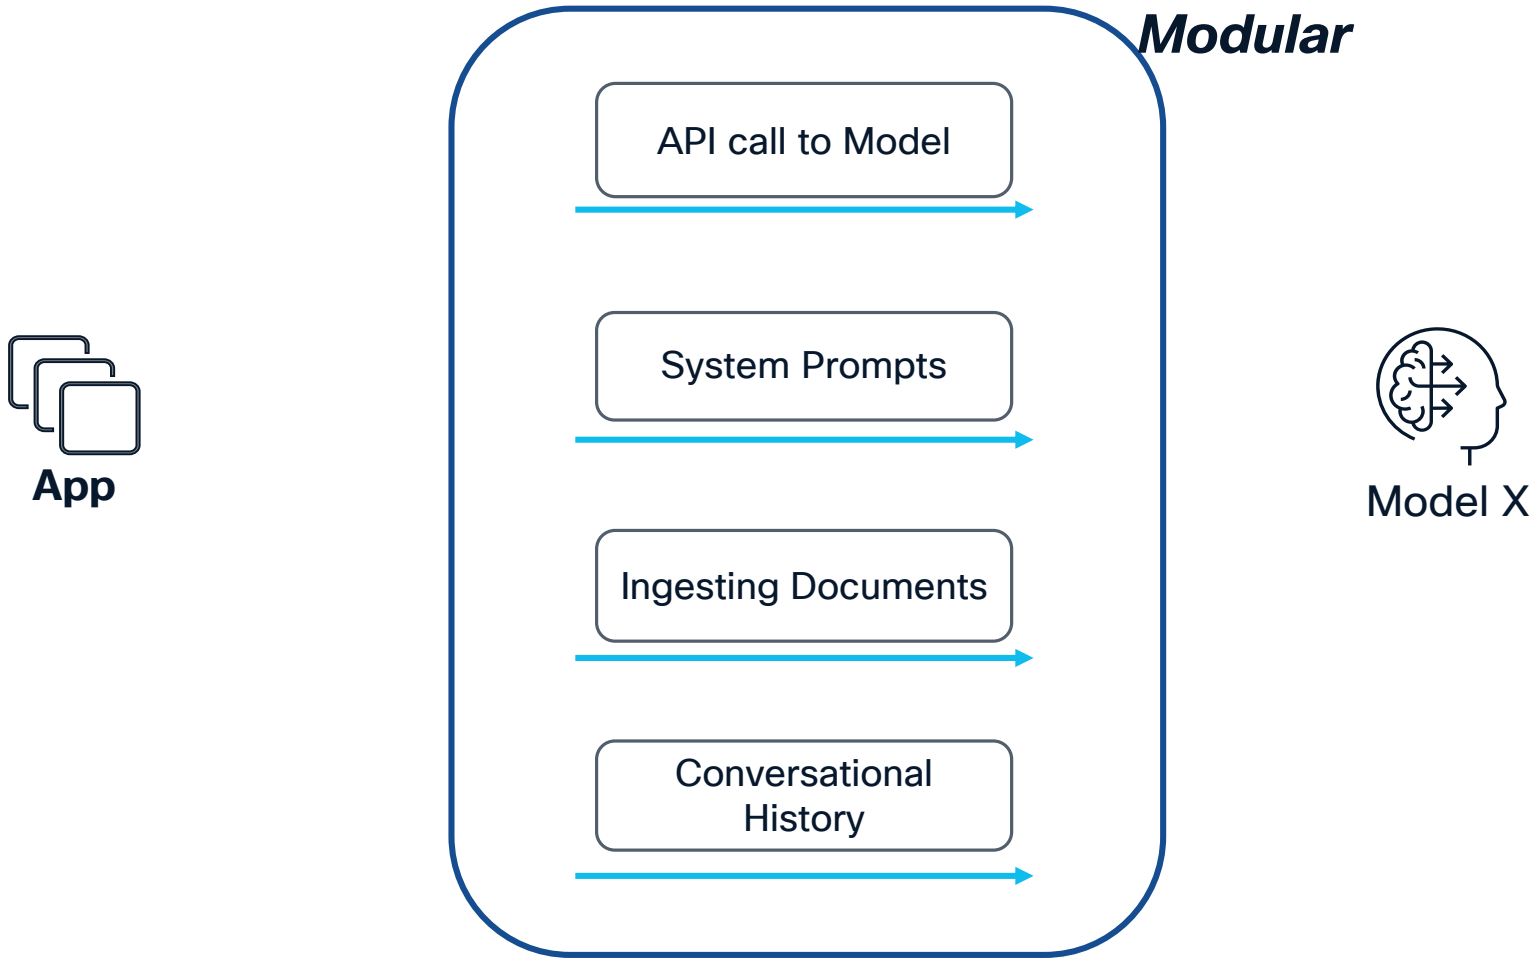

Integration Layer to the rescue -> Model Agnostic

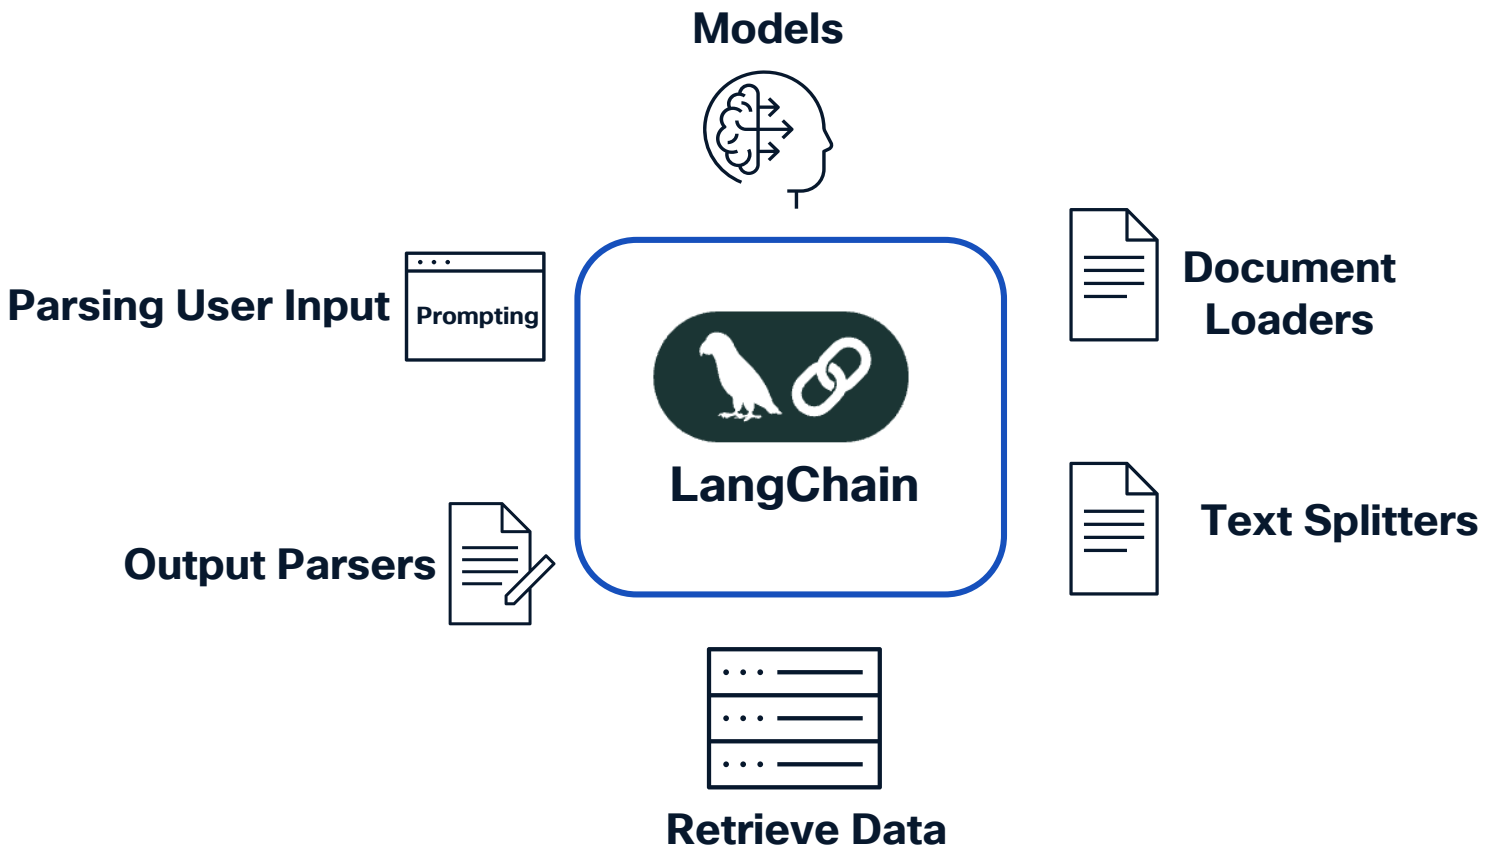

Linear Control Flow – a limiting factor

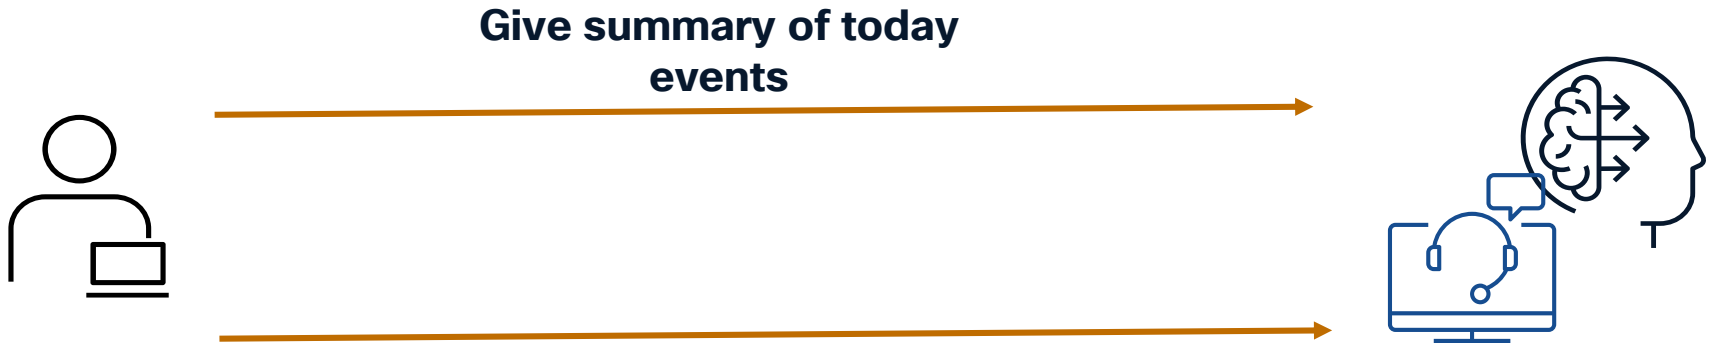

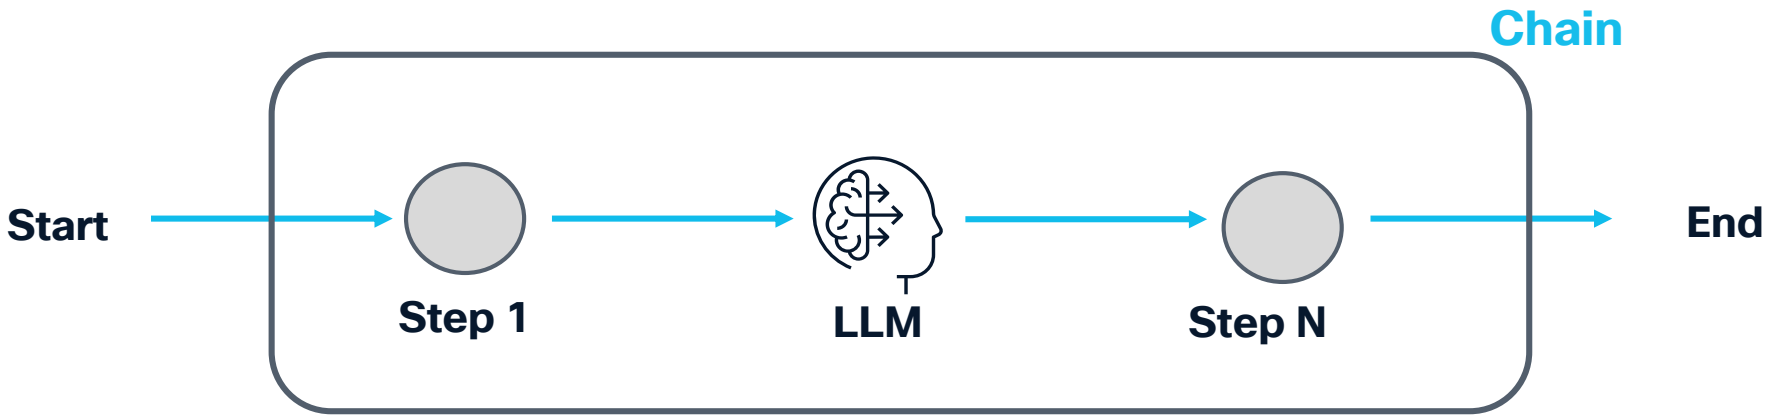

Non-Linear Control Flow with Agentic AI

Agentic

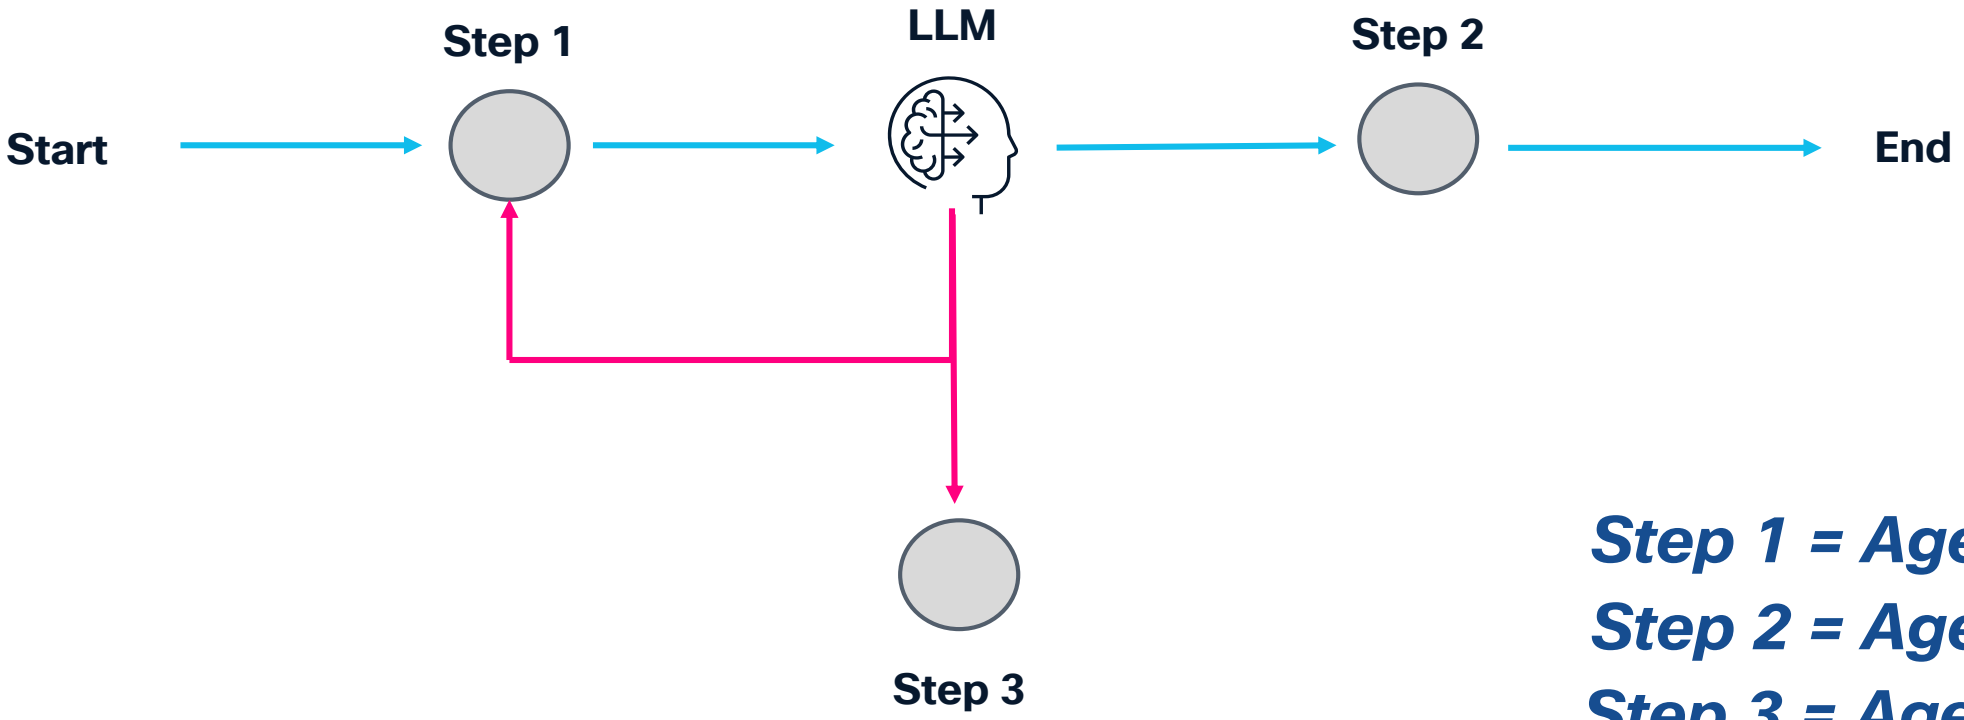

Step 1 = Agent 1
Step 2 = Agent 1
Step 3 = Agent 2

Building Blocks for Agentic AI

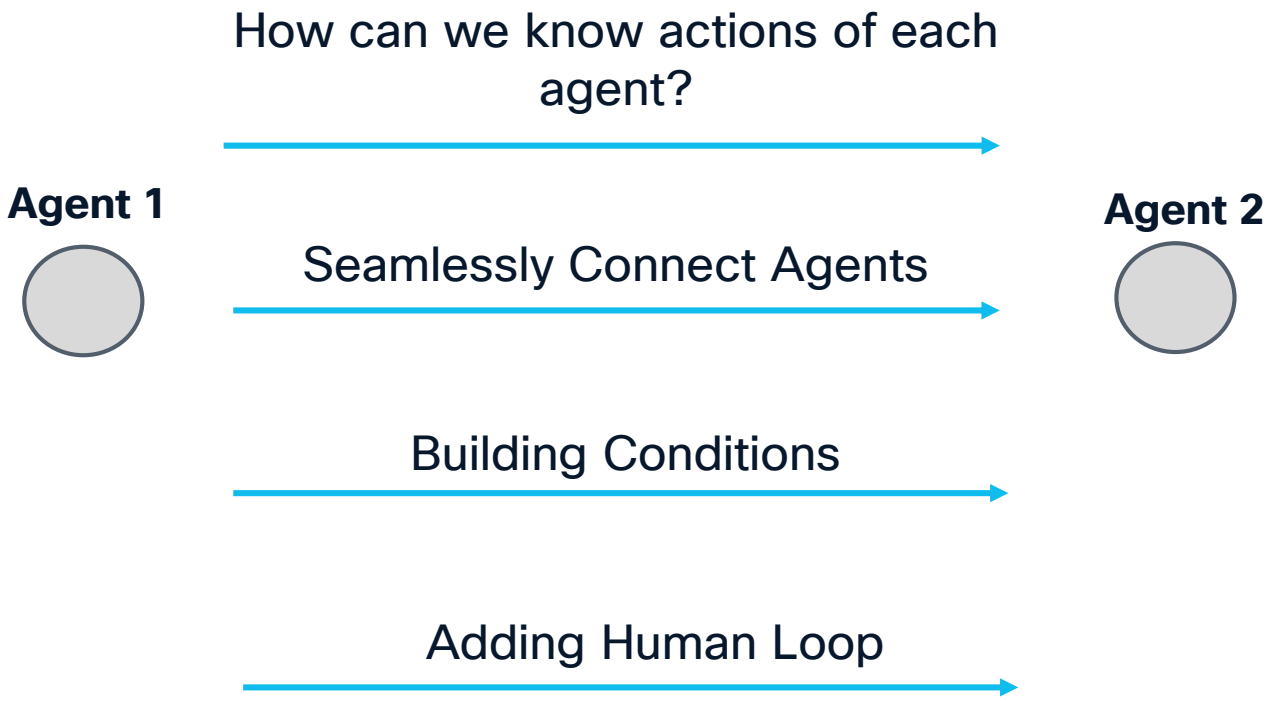

Orchestration Framework for Agentic AI

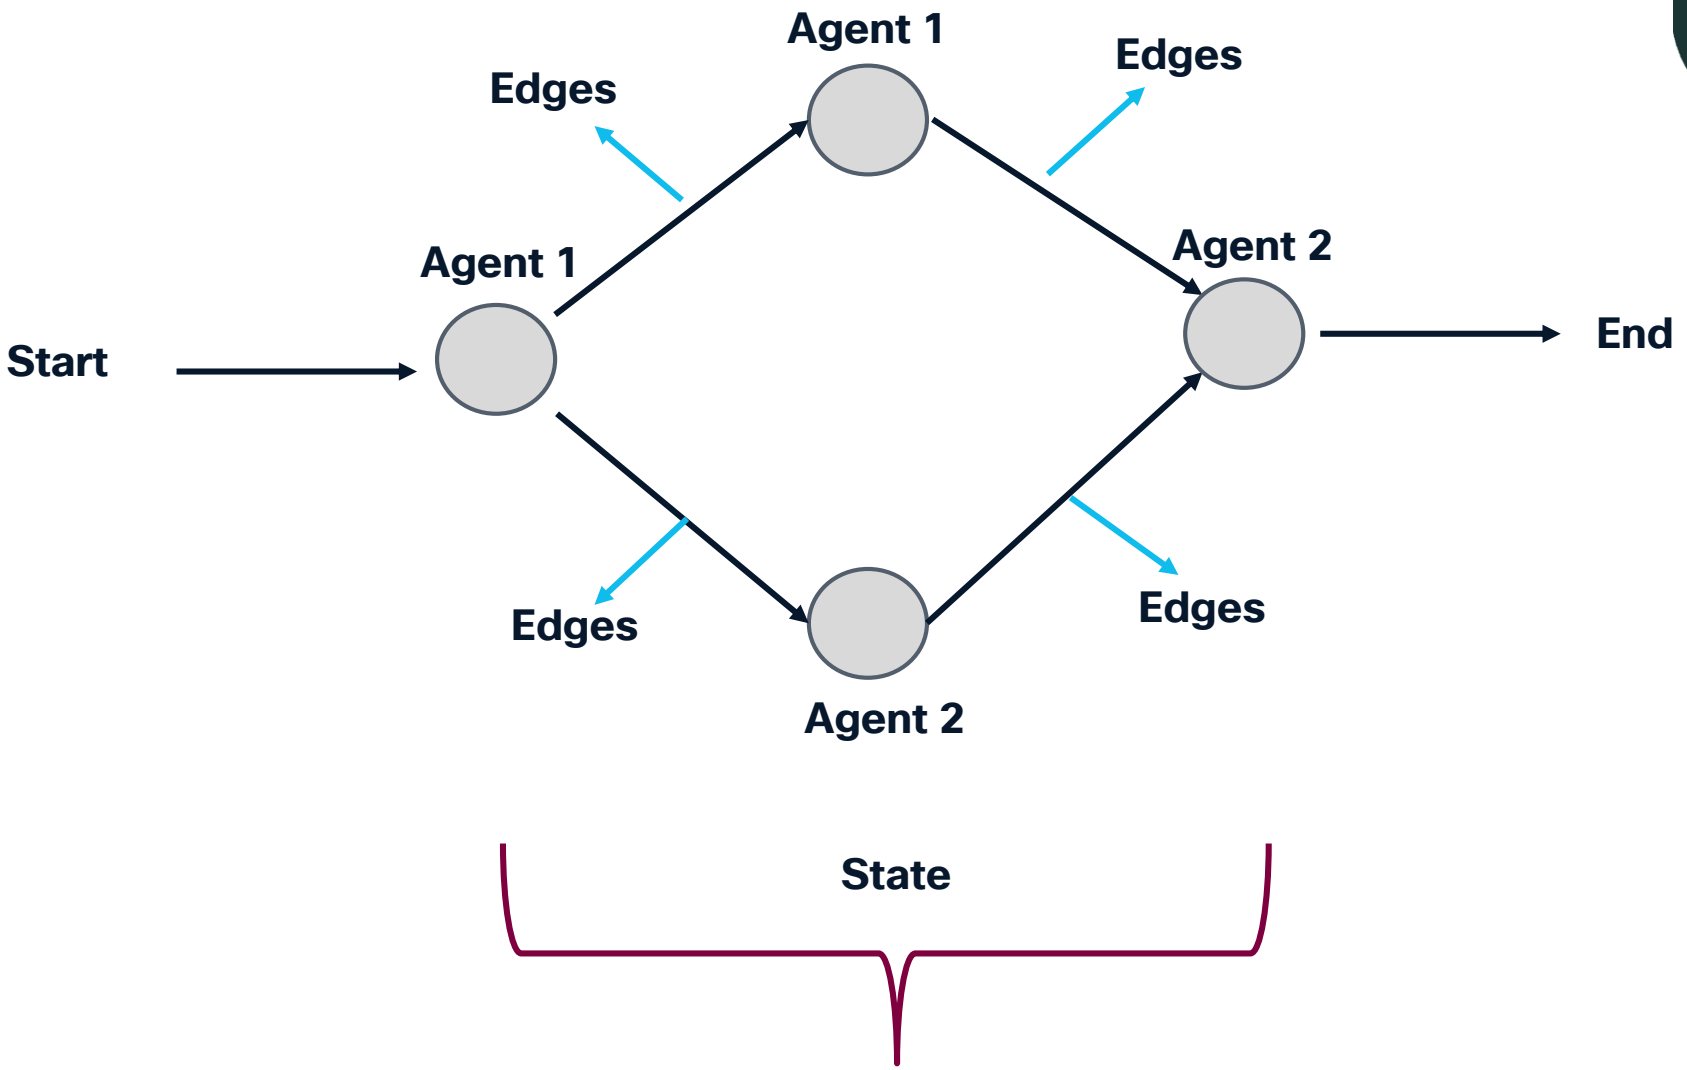

ModelOps

Retrieval Augmented Generation(RAG)

Definition:

- Allows LLM to query external data source to derive context

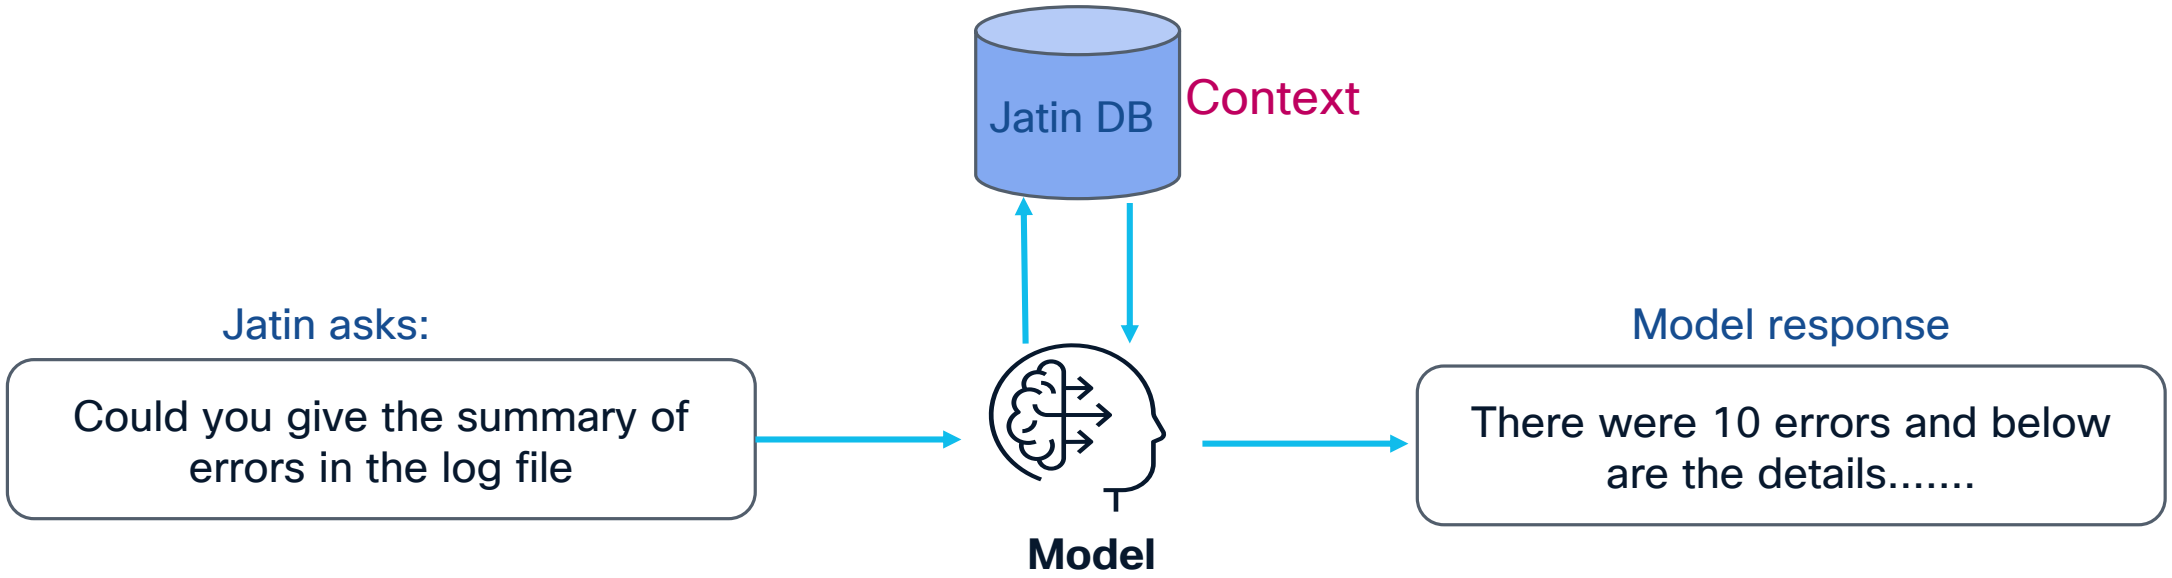

RAG Pipeline – Steps to make data LLM friendly

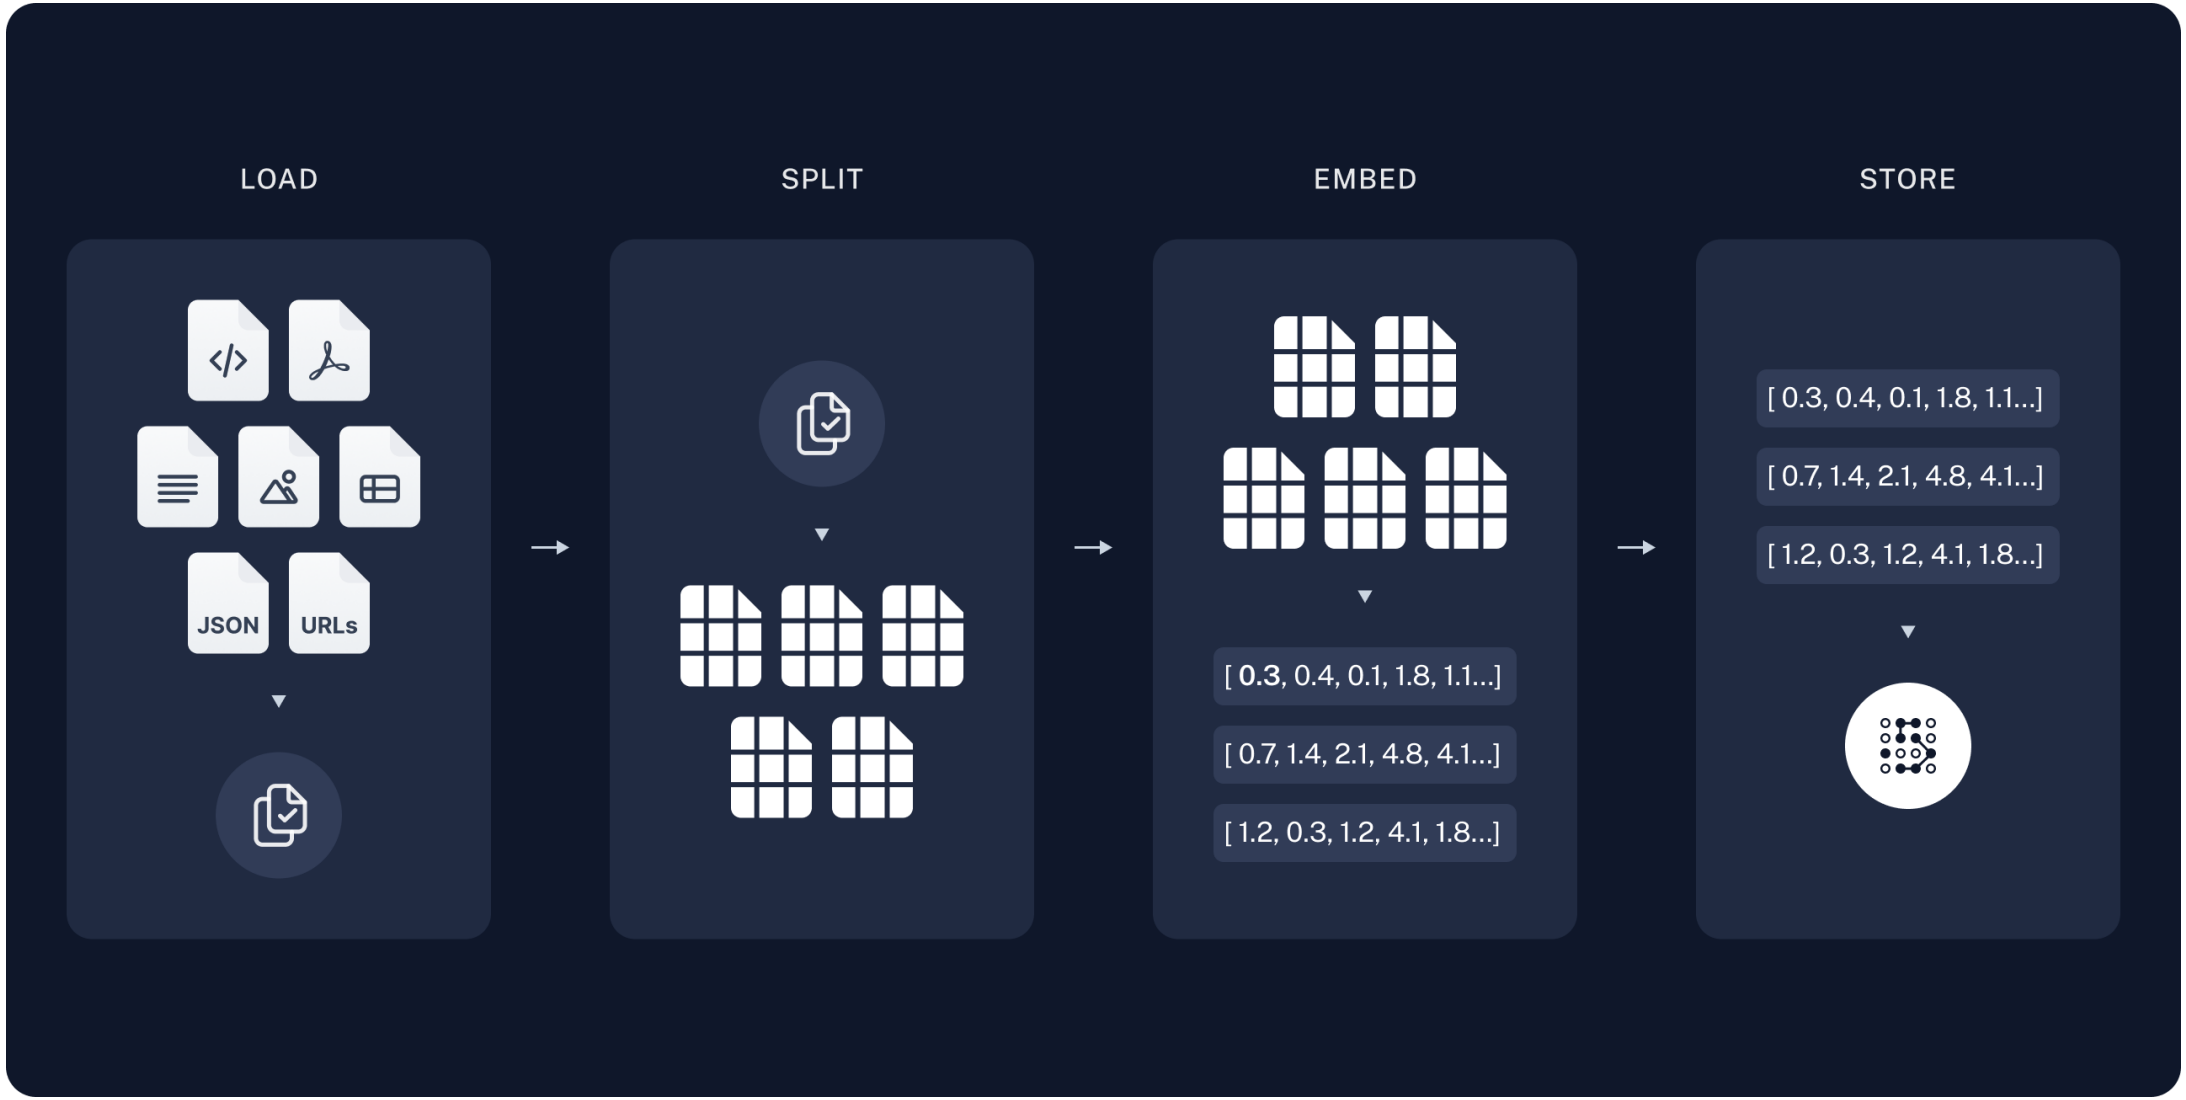

Jargon you may hear – for reference only

- **Embeddings :** Objects like text, audio, video and images are represented as "***Vectors***".
- **Embedding Models:** Specially trained ML models to represent embeddings in multi-dimensional space.
- **Vector Stores** : A database which is is specialized in storing the embeddings.
- **Tokens** : Units of text being processed

However RAG is becoming a Legacy option

Early 2024

32k tokens to 128k tokens

Currently

Rate limits

Rate limits ensure fair and reliable access to the API by placing specific caps on requests or tokens used within a given time period. Your usage tier determines how high these limits are set and automatically increases as you send more requests and spend more on the API.

TIER	RPM	TPM	BATCH QUEUE LIMIT
Free		Not supported	
Tier 1	500	30,000	90,000
Tier 2	5,000	450,000	1,350,000
Tier 3	5,000	800,000	50,000,000
Tier 4	10,000	2,000,000	200,000,000
Tier 5	10,000	30,000,000	5,000,000,000

Modern LLMs token limits make Fine Tuning easier

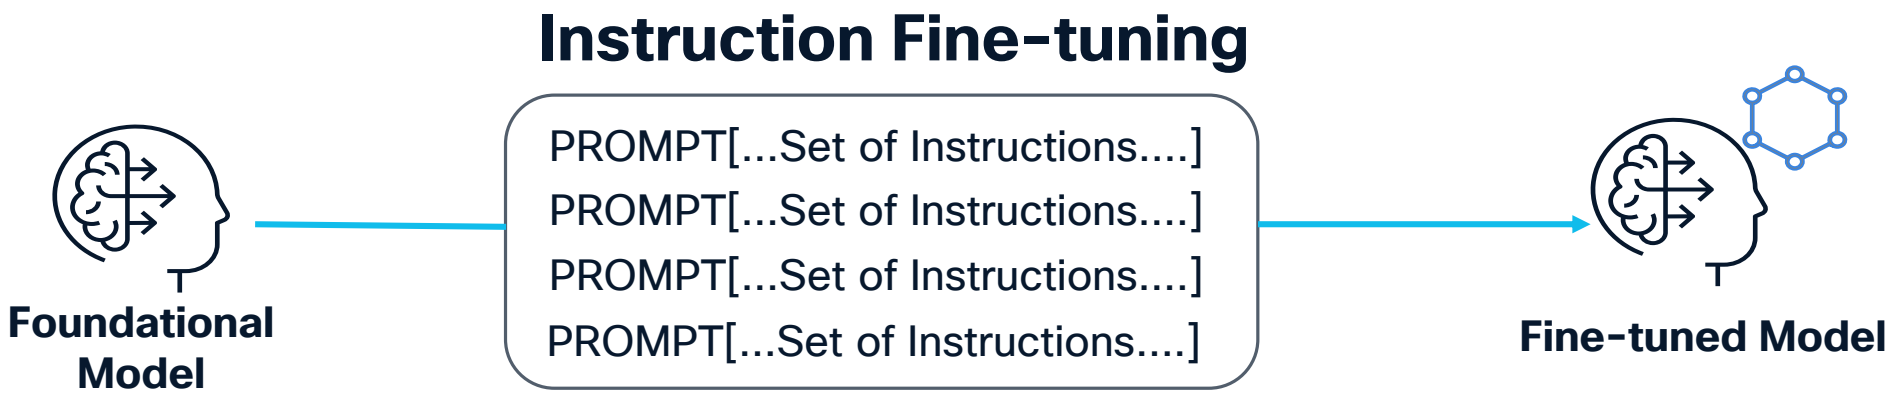

Using Prompts to Fine-Tune LLM

Fine Tuning for advanced use cases

Parameters : a , b , c

Weights : Y,X

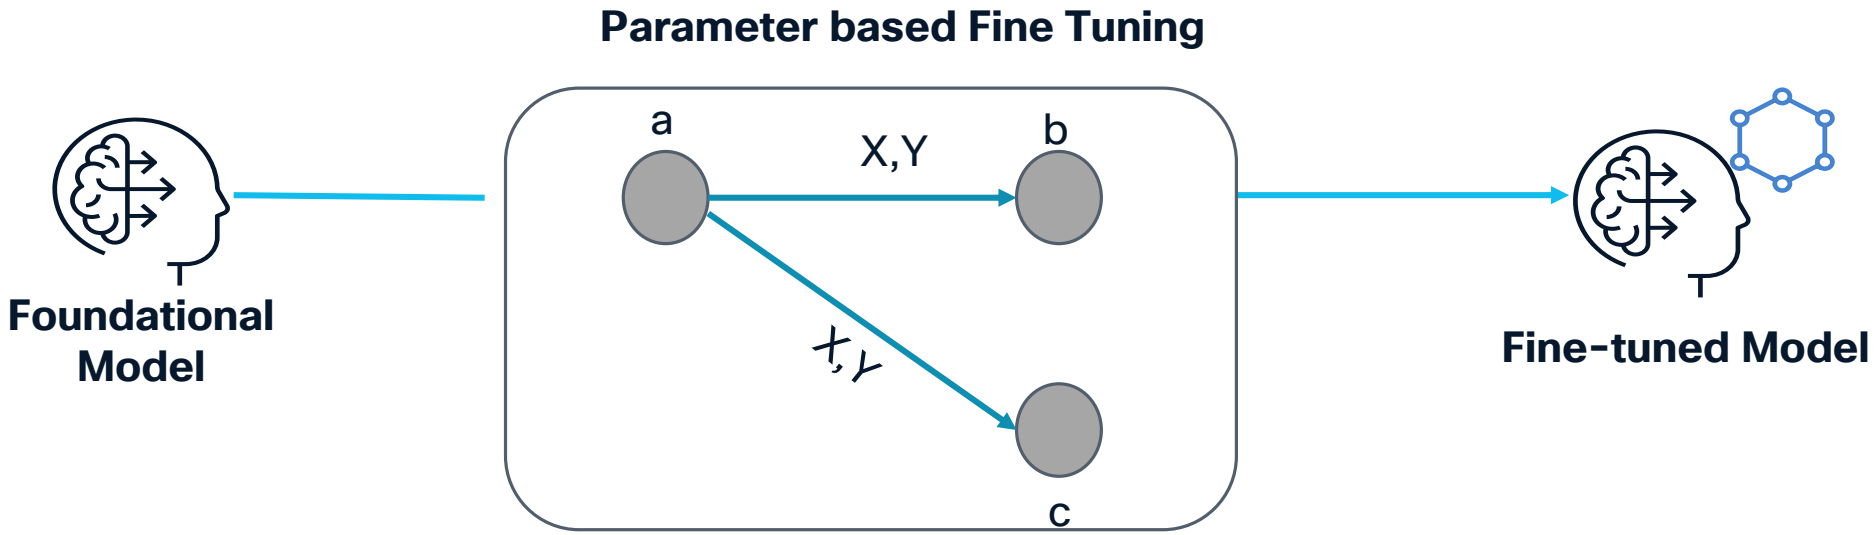

$Y = a^* X + b$

$Y = a^*X + c$

Model Deployment Example for Reference

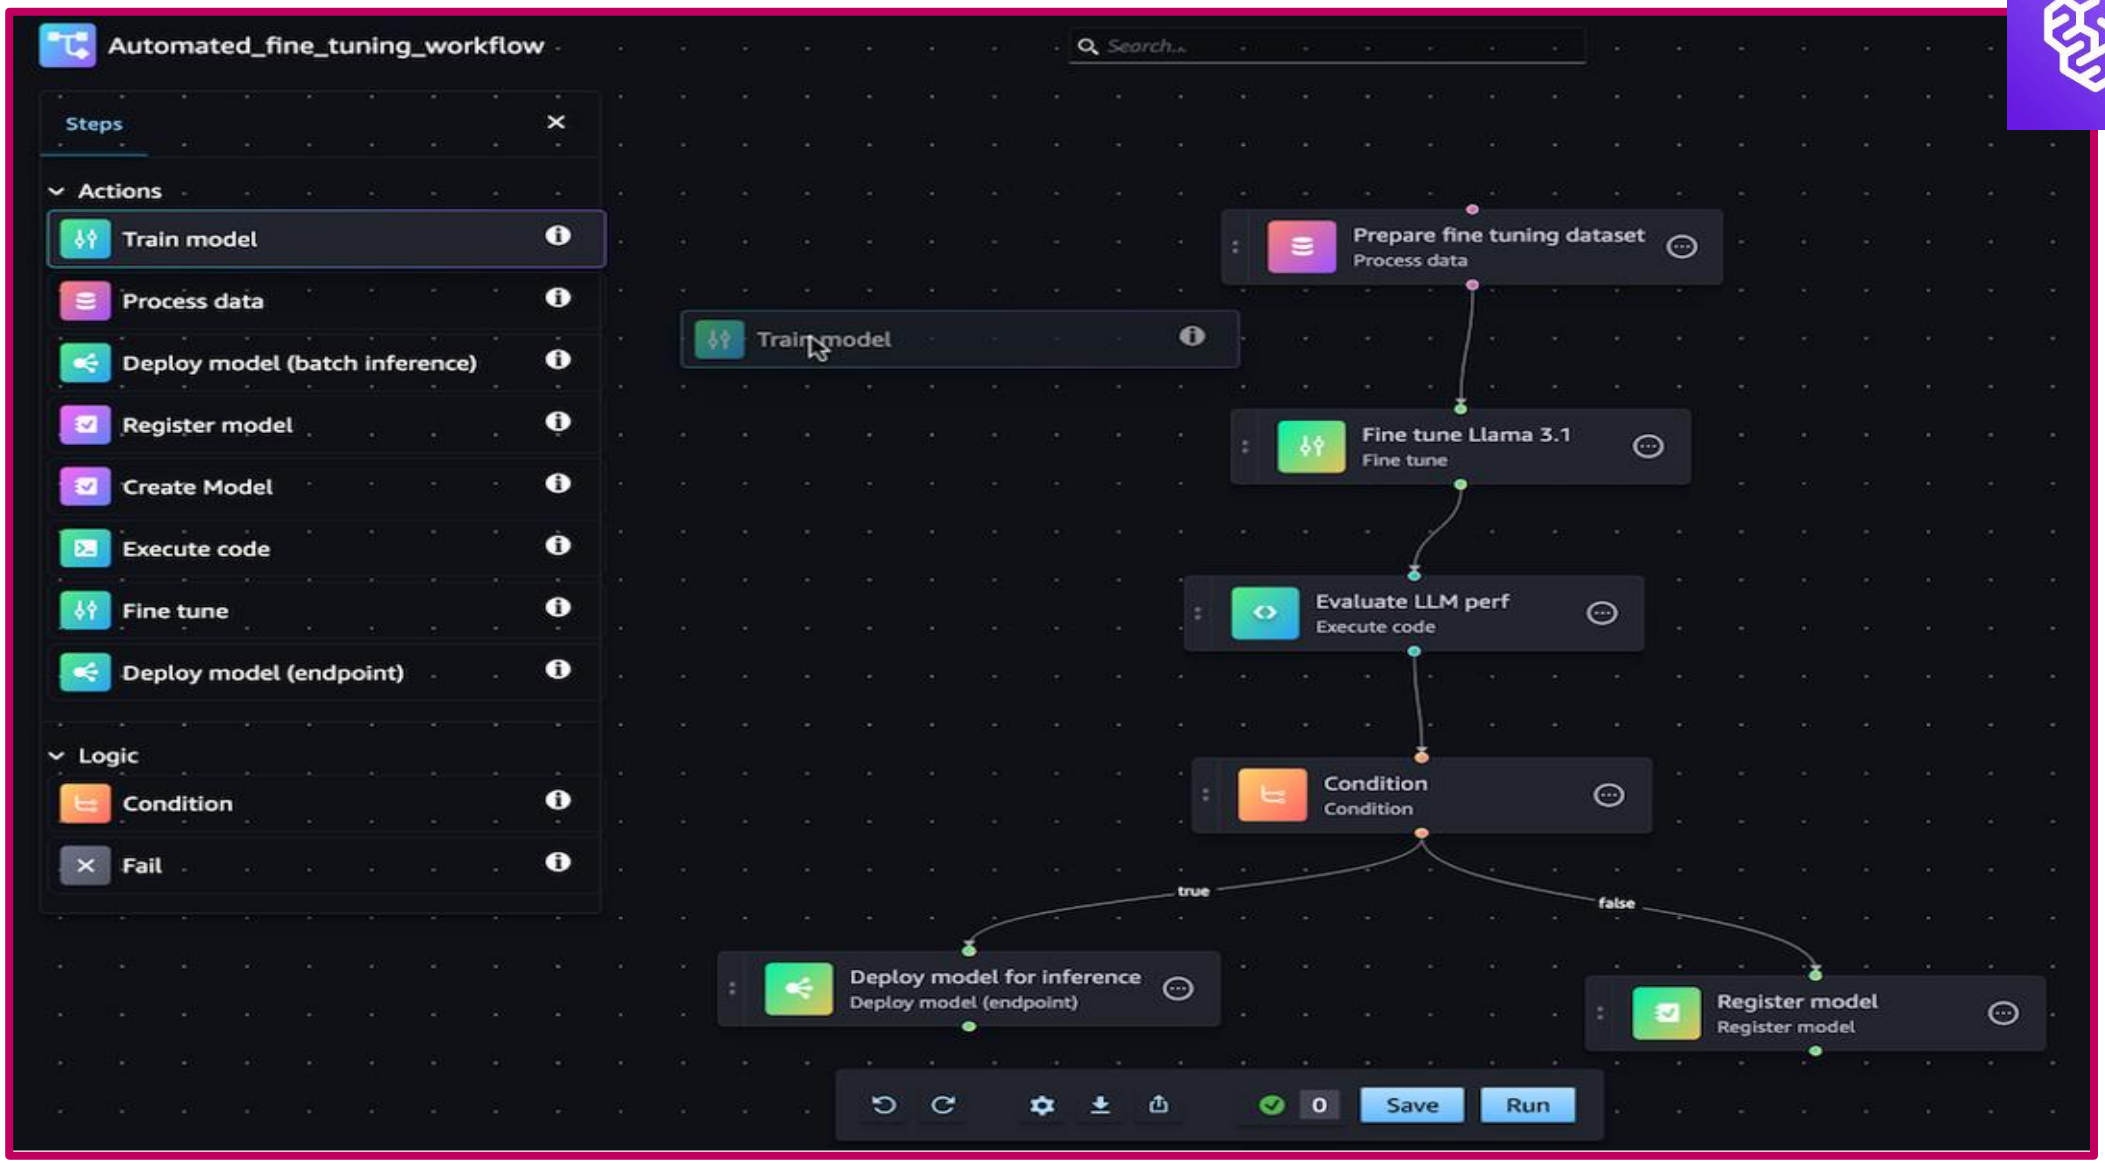

Our Sample LLM application

An example SecOps Use Case – Firewall Upgrades!

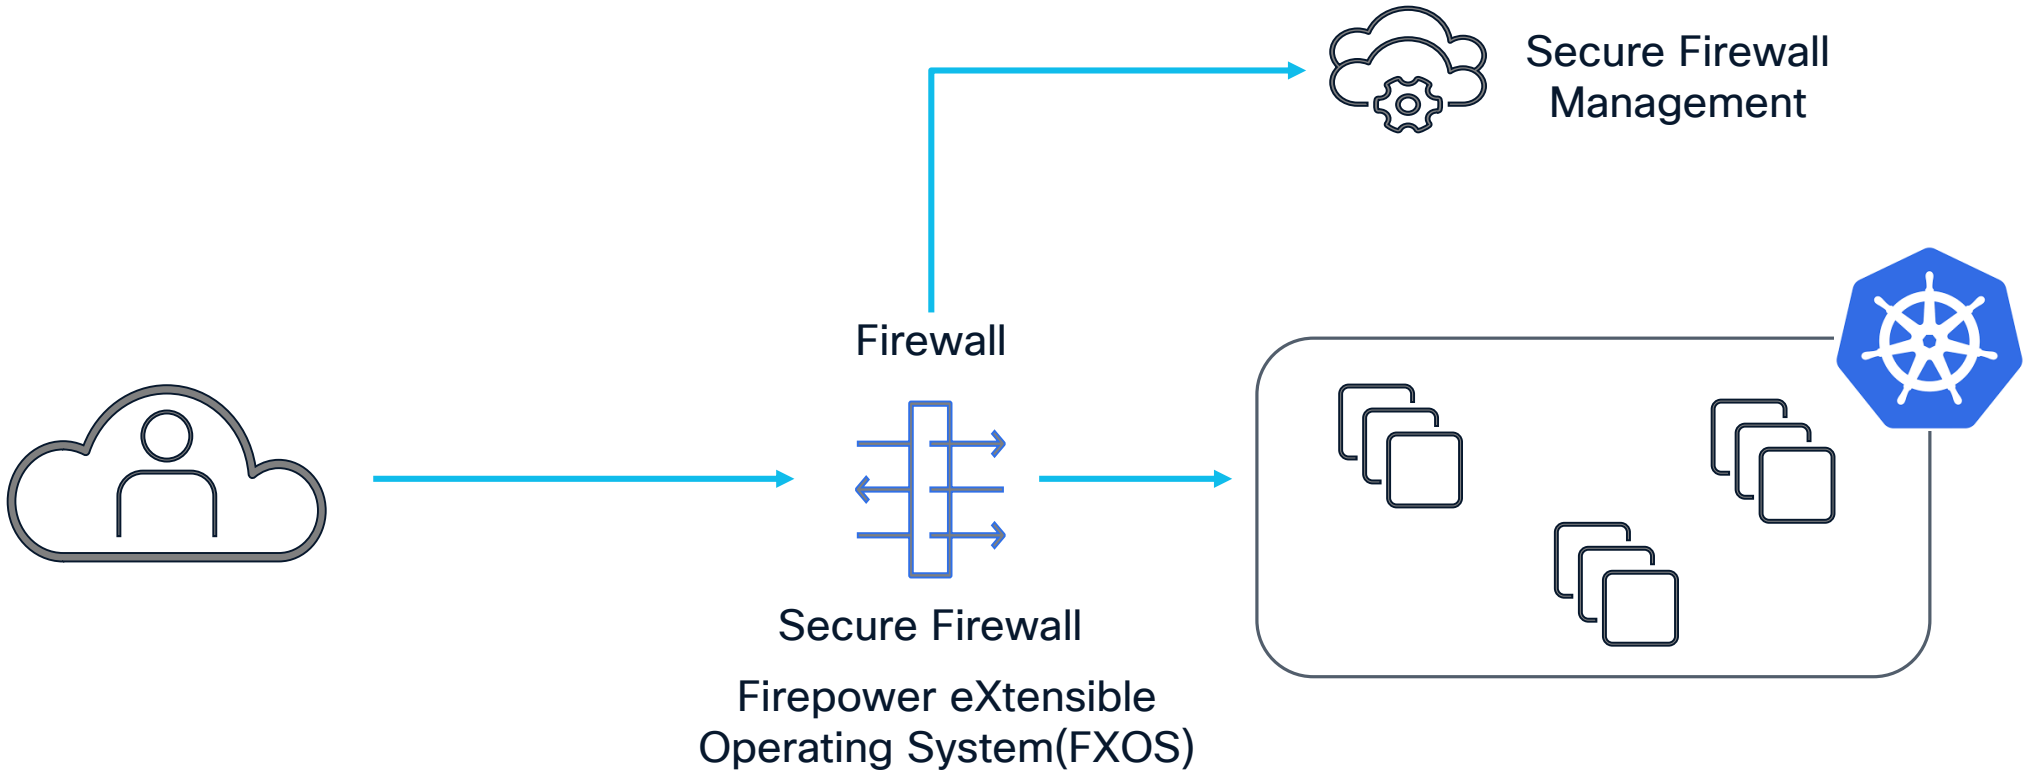

Sample LLM Application – Firewall Upgrade Assistant

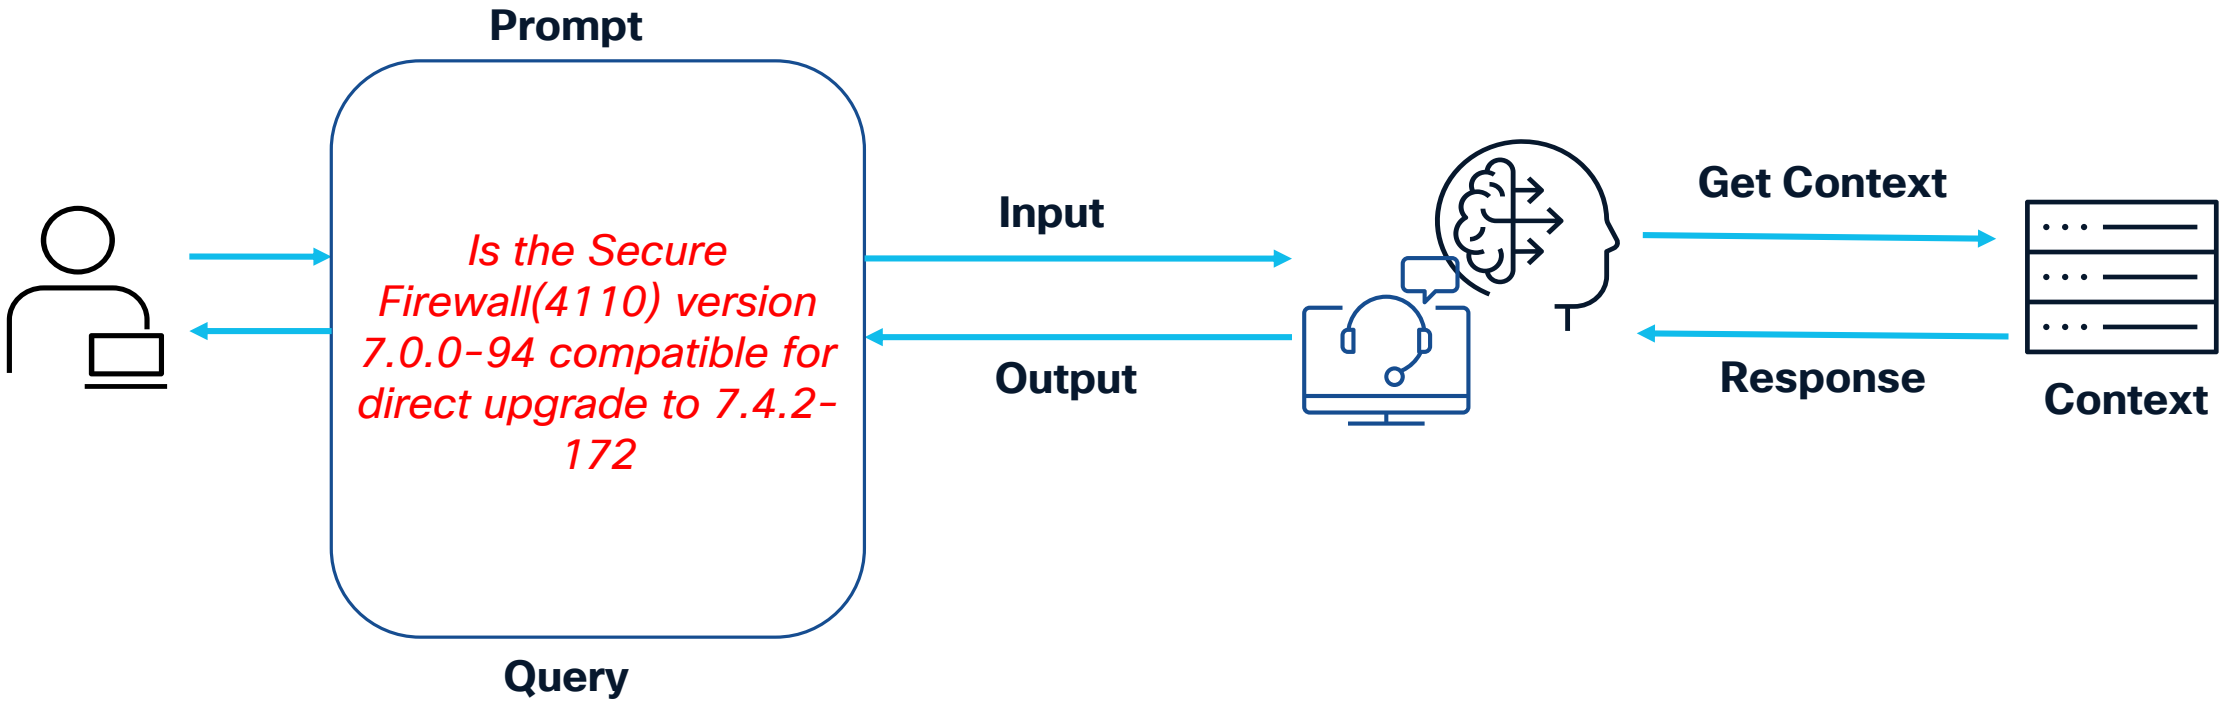

Upgrade assistant which answers user queries for different Cisco products

Simple Python based LLM app

LLM > main.py > ...

```python
from state import UpgradeState
from typing import Literal
from langchain_core.output_parsers import JsonOutputParser
from langchain.agents import create_react_agent, AgentExecutor, create_openai_functions_agent, OpenAIFunctionsA
from dotenv import load_dotenv
load_dotenv()



def compatibility_route(state)->Literal["tools", "end" ]:
    messages = state["messages"]
    last_message = messages[-1]
    if last_message.tool_calls:
        return "tools"
    else:
        return "end"
tools = [fmc_hardware_compatibility,fmc_hardware_bios_compatibility,fmc_virtual_300_public_cloud,fmc_virtual_pu
compatibility_check_tool_node = ToolNode(tools)

```

Simple Python based LLM app

LLM > main.py > ...

```
26  def compatibility_agent(state):
78      llm = ChatOpenAI(
79          model= "gpt-4o-mini",
80          temperature=0
81      )
82
83      compatibility_with_tools = llm.bind_tools(tools)
84
85      chat_prompt_template = chat_prompt_template.partial(compatibility_prompt=compatibility_prompt,tools=tools)
86
87      chat_prompt_template = chat_prompt_template | compatibility_with_tools
88      #print("chat prompt template\n",chat_prompt_template)
89      chat_response = chat_prompt_template.invoke(state)
90      resp_dict = {"messages": [chat_response]}
```

Knowledge Sources

Management Center Compatibility Guide → https://www.cisco.com/c/en/us/td/docs/security/secure-firewall/compatibility/management-center-compatibility.html

Secure Firewall Threat Defense Compatibility Guide → https://www.cisco.com/c/en/us/td/docs/security/secure-firewall/compatibility/threat-defense-compatibility.html

Firepower Compatibility Guide → https://www.cisco.com/c/en/us/td/docs/security/firepower/fxos/compatibility/fxos-compatibility.html

ASA Compatibility Guide → https://www.cisco.com/c/en/us/td/docs/security/asa/compatibility/asamatrx.html

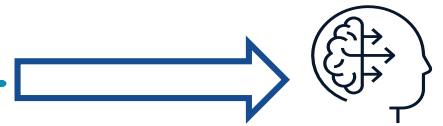

Model

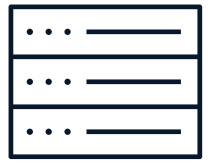

Context

Knowledge sources in code

```
LLM > main.py > compatibility_agent
16  def compatibility_route(state)->Literal["tools", "end" ]:
18      last_message = messages[-1]
19      if last_message.tool_calls:
20          return "tools"
21      else:
22          return "end"
23  (variable) compatibility_check_tool_node: ToolNode
24  compatibility_check_tool_node = ToolNode(tools)
25
26  def compatibility_agent(state):
27
28
29
30
31      tools = [fmc_hardware_compatibility,fmc_hardware_bios_compatibility,fmc_virtual_300_public_cloud,fmc_virtual_public_cloud,fmc_virtual
32
33
34      COMP_URL = [
35          "https://www.cisco.com/c/en/us/td/docs/security/secure-firewall/compatibility/threat-defense-compatibility.html",
36          "https://www.cisco.com/c/en/us/td/docs/security/firepower/fxos/compatibility/fxos-compatibility.html",
37          "https://www.cisco.com/c/en/us/td/docs/security/secure-firewall/compatibility/management-center-compatibility.html",
38          "https://www.cisco.com/c/en/us/td/docs/security/asa/compatibility/asamatrx.html"
39
40      ]
41
42
43      #correct template1
44
45      compatibility_prompt = SystemMessage(f"""You are a TAC expert and your primary role is to solve compatibility queries for different C
46                                           reference different tables which are present in a variable 'table_name' in tools. Do not generat
47                                           You need to retrieve tables from the tools and the tables would give the information in terms of
48                                           Read the table information and answer if the product/versions are compatible or not. If not prov
49                                           Select the URL defined in list COMP_URL depending on the tool prefix.
```

Functions

Table 9. On-Prem Management Center- Device Compatibility

Management Center Version	Oldest Device Version You Can Manage
7.7	7.2
7.6	7.1
7.4 Last support for NGIPS device management.	7.0
7.3	6.7
7.2	6.6
7.1	6.5
7.0	6.4
6.7	6.3
6.6	6.2.3
6.5	6.2.3
6.4	6.1
6.3	6.1
6.2.3	6.1
6.2.2	6.1
6.2.1	6.1
6.2	6.1
6.1	5.4.0.2/5.4.1.1
6.0.1	5.4.0.2/5.4.1.1
6.0	5.4.0.2/5.4.1.1
5.4.1	5.4.1 for ASA FirePOWER on the ASA-5506-X series, ASA5508-X, and ASA5516-X. 5.3.1 for ASA FirePOWER on the ASA5512-X, ASA5515-X, ASA5525-X, ASA5545-X, ASA5555-X, and ASA-5585-X series. 5.3.0 for Firepower 7000/8000 series and legacy devices.

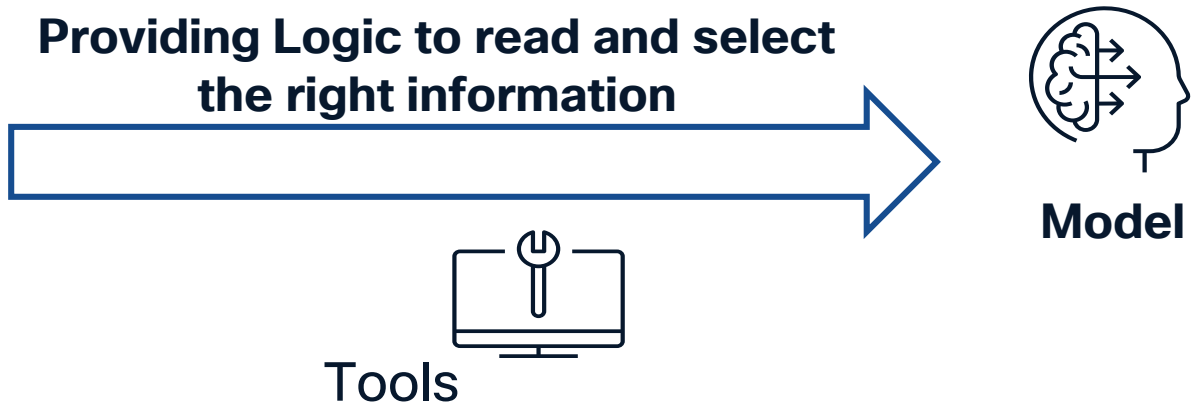

Knowledge sources in code

```
LLM > main.py > compatibility_agent
16  def compatibility_route(state)->Literal["tools", "end" ]:
18      last_message = messages[-1]
19      if last_message.tool_calls:
20          return "tools"
21      else:
22          return "end"
23  (variable) compatibility_check_tool_node: ToolNode
24  compatibility_check_tool_node = ToolNode(tools)
25
26  def compatibility_agent(state):
27
28
29
30
31      tools = [fmc_hardware_compatibility,fmc_hardware_bios_compatibility,fmc_virtual_300_public_cloud,fmc_virtual_public_cloud,fmc_virtual
32
33
34      COMP_URL = [
35          "https://www.cisco.com/c/en/us/td/docs/security/secure-firewall/compatibility/threat-defense-compatibility.html",
36          "https://www.cisco.com/c/en/us/td/docs/security/firepower/fxos/compatibility/fxos-compatibility.html",
37          "https://www.cisco.com/c/en/us/td/docs/security/secure-firewall/compatibility/management-center-compatibility.html",
38          "https://www.cisco.com/c/en/us/td/docs/security/asa/compatibility/asamatrx.html"
39
40      ]
41
42
43      #correct template1
44
45      compatibility_prompt = SystemMessage(f"""You are a TAC expert and your primary role is to solve compatibility queries for different C
46                                           reference different tables which are present in a variable 'table_name' in tools. Do not generat
47                                           You need to retrieve tables from the tools and the tables would give the information in terms of
48                                           Read the table information and answer if the product/versions are compatible or not. If not prov
49                                           Select the URL defined in list COMP_URL depending on the tool prefix.
```

Tools – function that converts tables to JSON

LLM > tools.py > ...

```python
@tool
def ftd_3100_4200_compatibility(URL):
    """
        This function gives the information on compatibility for 3100 and 4200 ftd devices."""
    table_name = "Secure Firewall 3100/4200 Series Application Mode Compatibility"
    heading = None
    table_context = table_ref_id(URL)
    return table_context

@tool
def ftd_multi_instance_compatibility(URL):
    """
        This function gives the information on multi-instance mode for 3100/4200 ftd devices."""
    table_name = "Secure Firewall 3100/4200 Series Multi-Instance Mode Compatibility"
    heading = None
    table_context = table_ref_id(URL,table_name,heading)
    return table_context

@tool
def ftd_4100_9300_compatibility(URL):
    """
        This function gives the information on ftd and fxos compatibility for 4100/9300 ftd devices"""
    table_name = "Firepower 4100/9300 Compatibility"
    heading = None
    table_context = table_ref_id(URL,table_name,heading)
    return table_context
```

Fine Tuning with Prompt Engineering

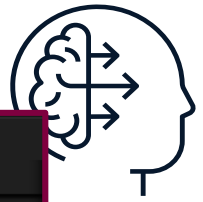

LLM > main.py > compatibility_agent

```
26  def compatibility_agent(state):
45      compatibility_prompt = SystemMessage(f"""You are a TAC expert and your primary role is to solve compatibility queries for different C
46                                           reference different tables which are present in a variable 'table_name' in tools. Do not generat
47                                           You need to retrieve tables from the tools and the tables would give the information in terms of
48                                           Read the table information and answer if the product/versions are compatible or not. If not prov
49                                           Select the URL defined in list COMP_URL depending on the tool prefix.
50
51          1) If the tool name starts with "fmc" kindly use the URL ending in "management-center-compatibility.html"
52          2) If the tool name starts with "ftd" kindly use the URL ending in "threat-defense-compatibility.html"
53          3) If the tool name starts with "asa" kindly use the URL ending in "asamatrix.html"
54          4) If the tool name starts with "fxos" kindly use the URL ending in "fxos-compatibility.html"
55          Below are few acronyms which you should be aware of
56          FMC : Secure Firewall Management Center
57          FTD : Secure Firewall Thread Defense
58          FXOS: Firepower eXtensible Operating System
59          ASA : Adaptive Security Appliance
60          The query input would be given by the user
61          The tools would give the context. Each tool has a description and try to map the query to the description provided. These would h
62          choosing the correct table for answering the query.
63          The list of URL is defined in COMP_URL={COMP_URL}
64          Return the response if the products are compatible to direct or step upgrade required.
65
66      """)
67      chat_prompt_template= ChatPromptTemplate.from_messages(
68          [
69              (
70                  "system",
71                  "You are intelligent assistant which resolves compatible queries."
72                  "Your response should be in 100 words."
73                  "You have access to the following tools {tools}\n {compatibility_prompt}"
74              ),
75              MessagesPlaceholder("messages")
```

Observability

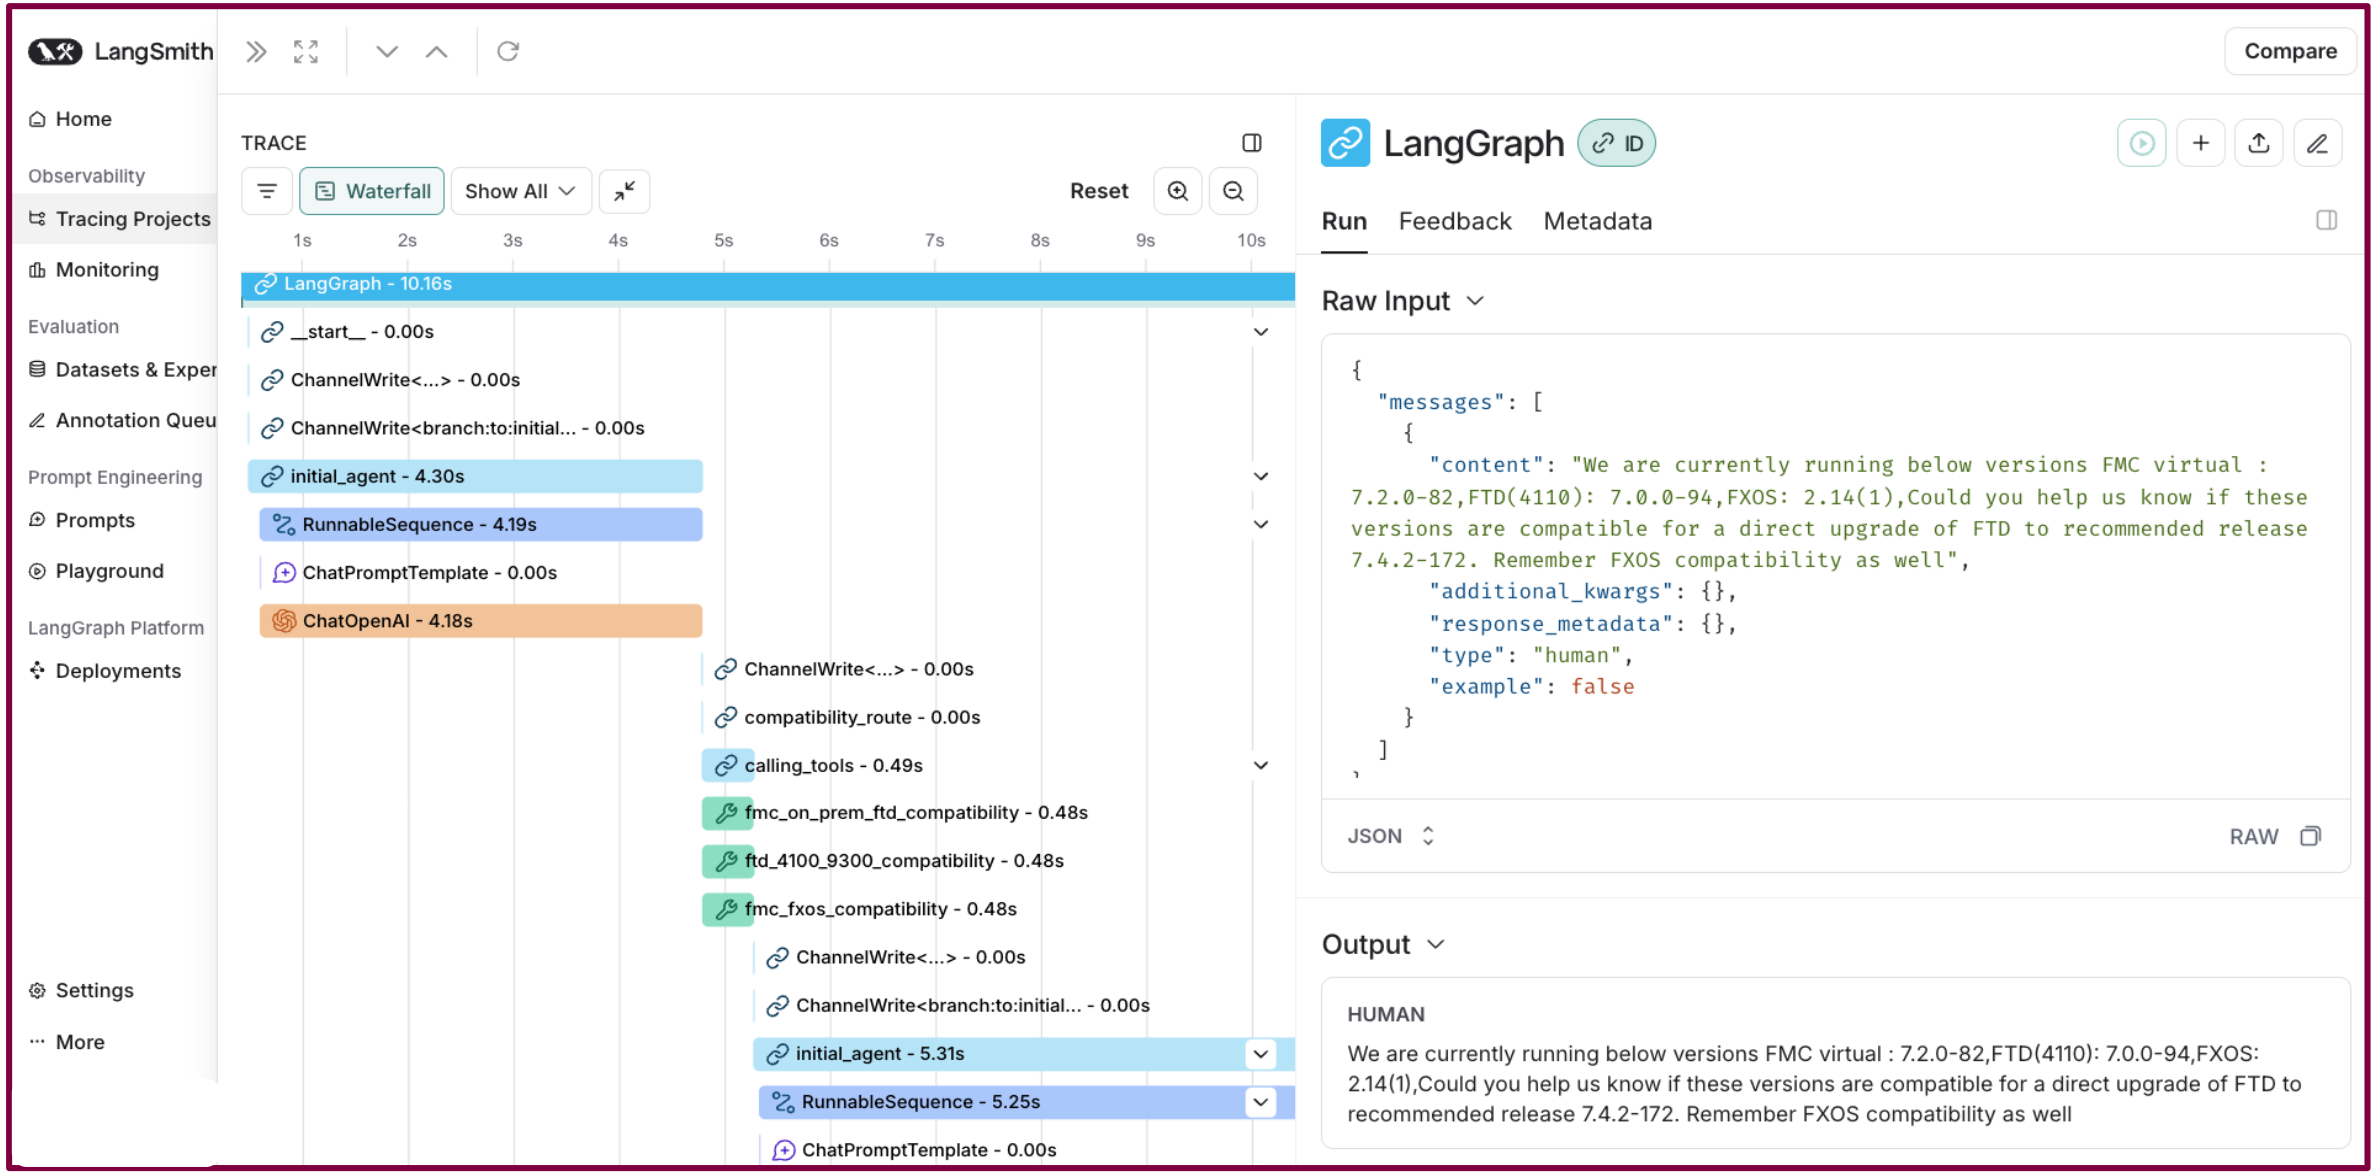

Putting together our LLM Application

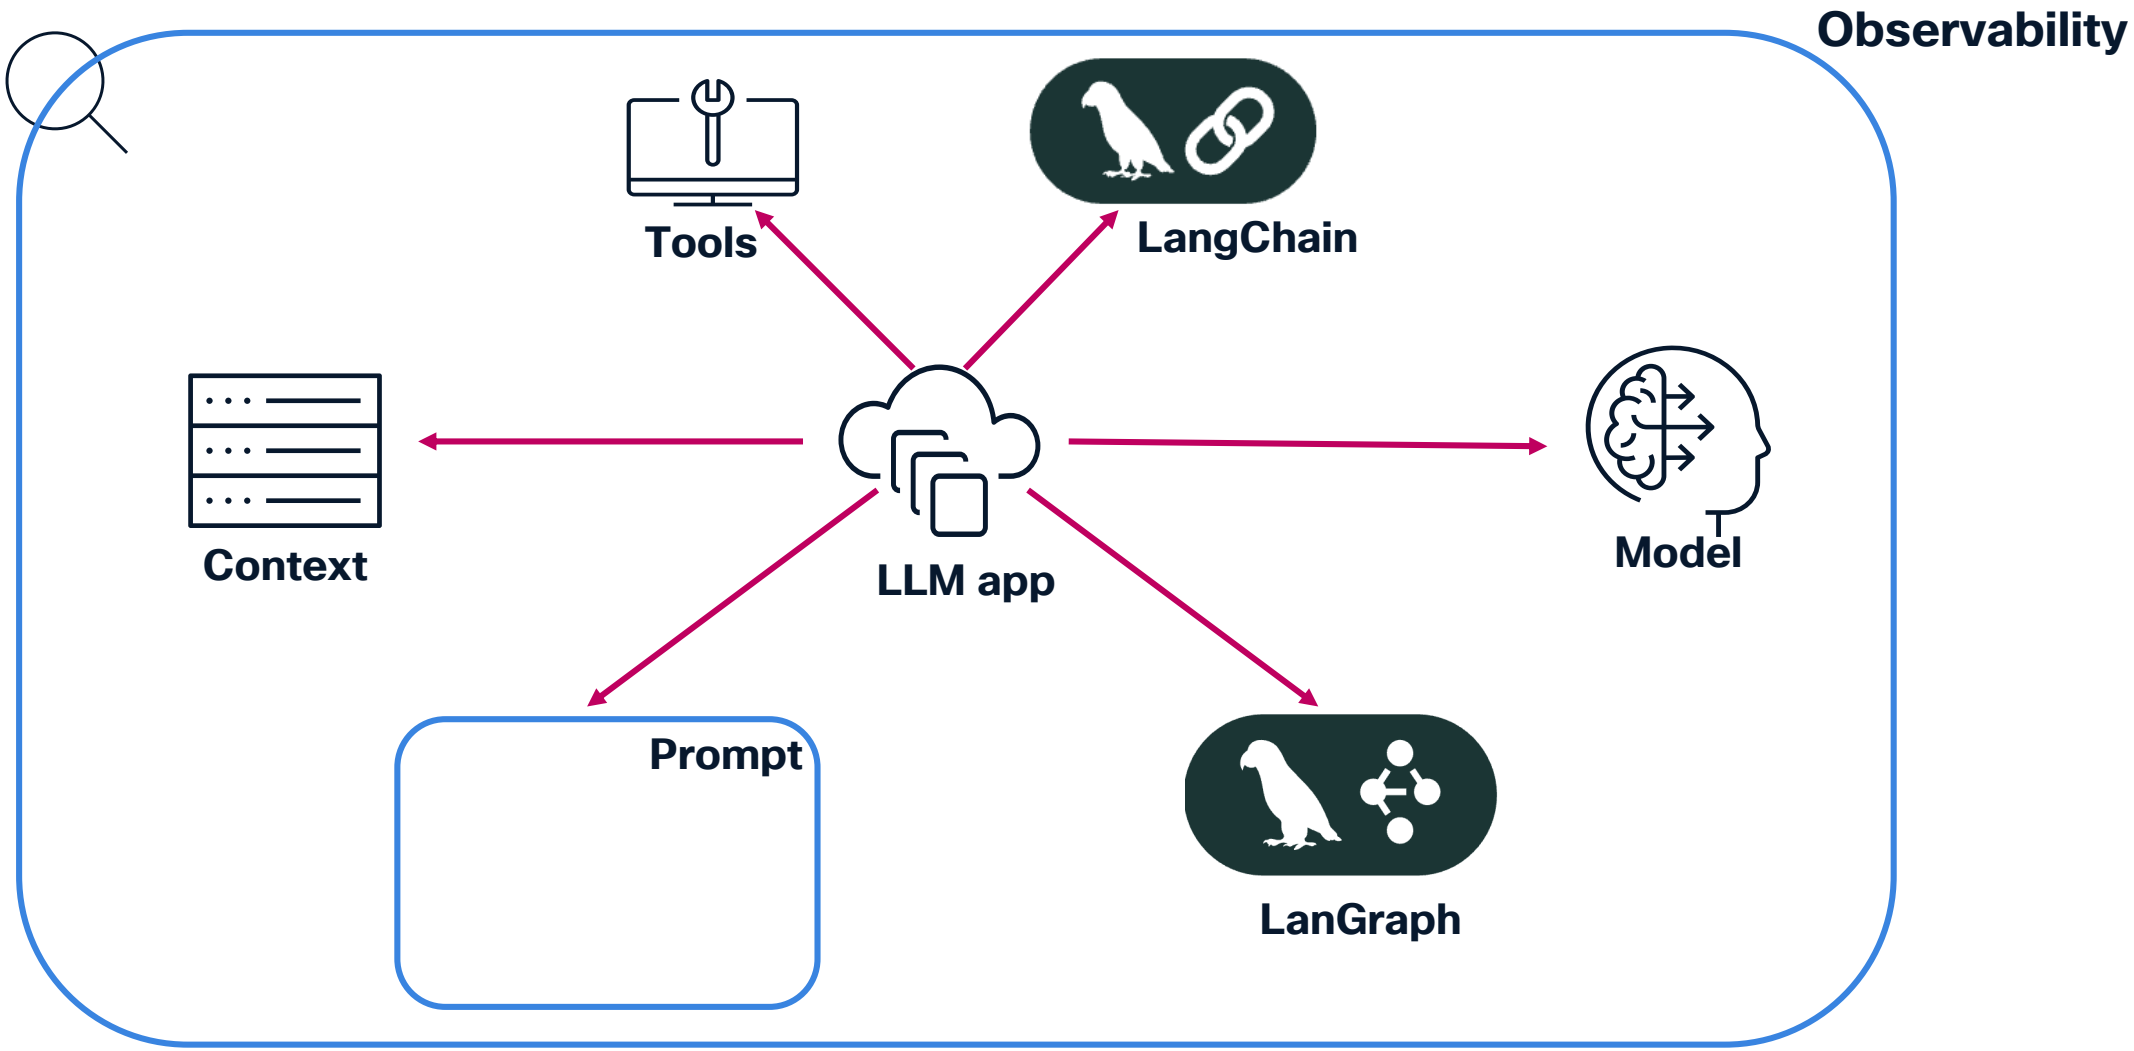

Quick demo of our final app

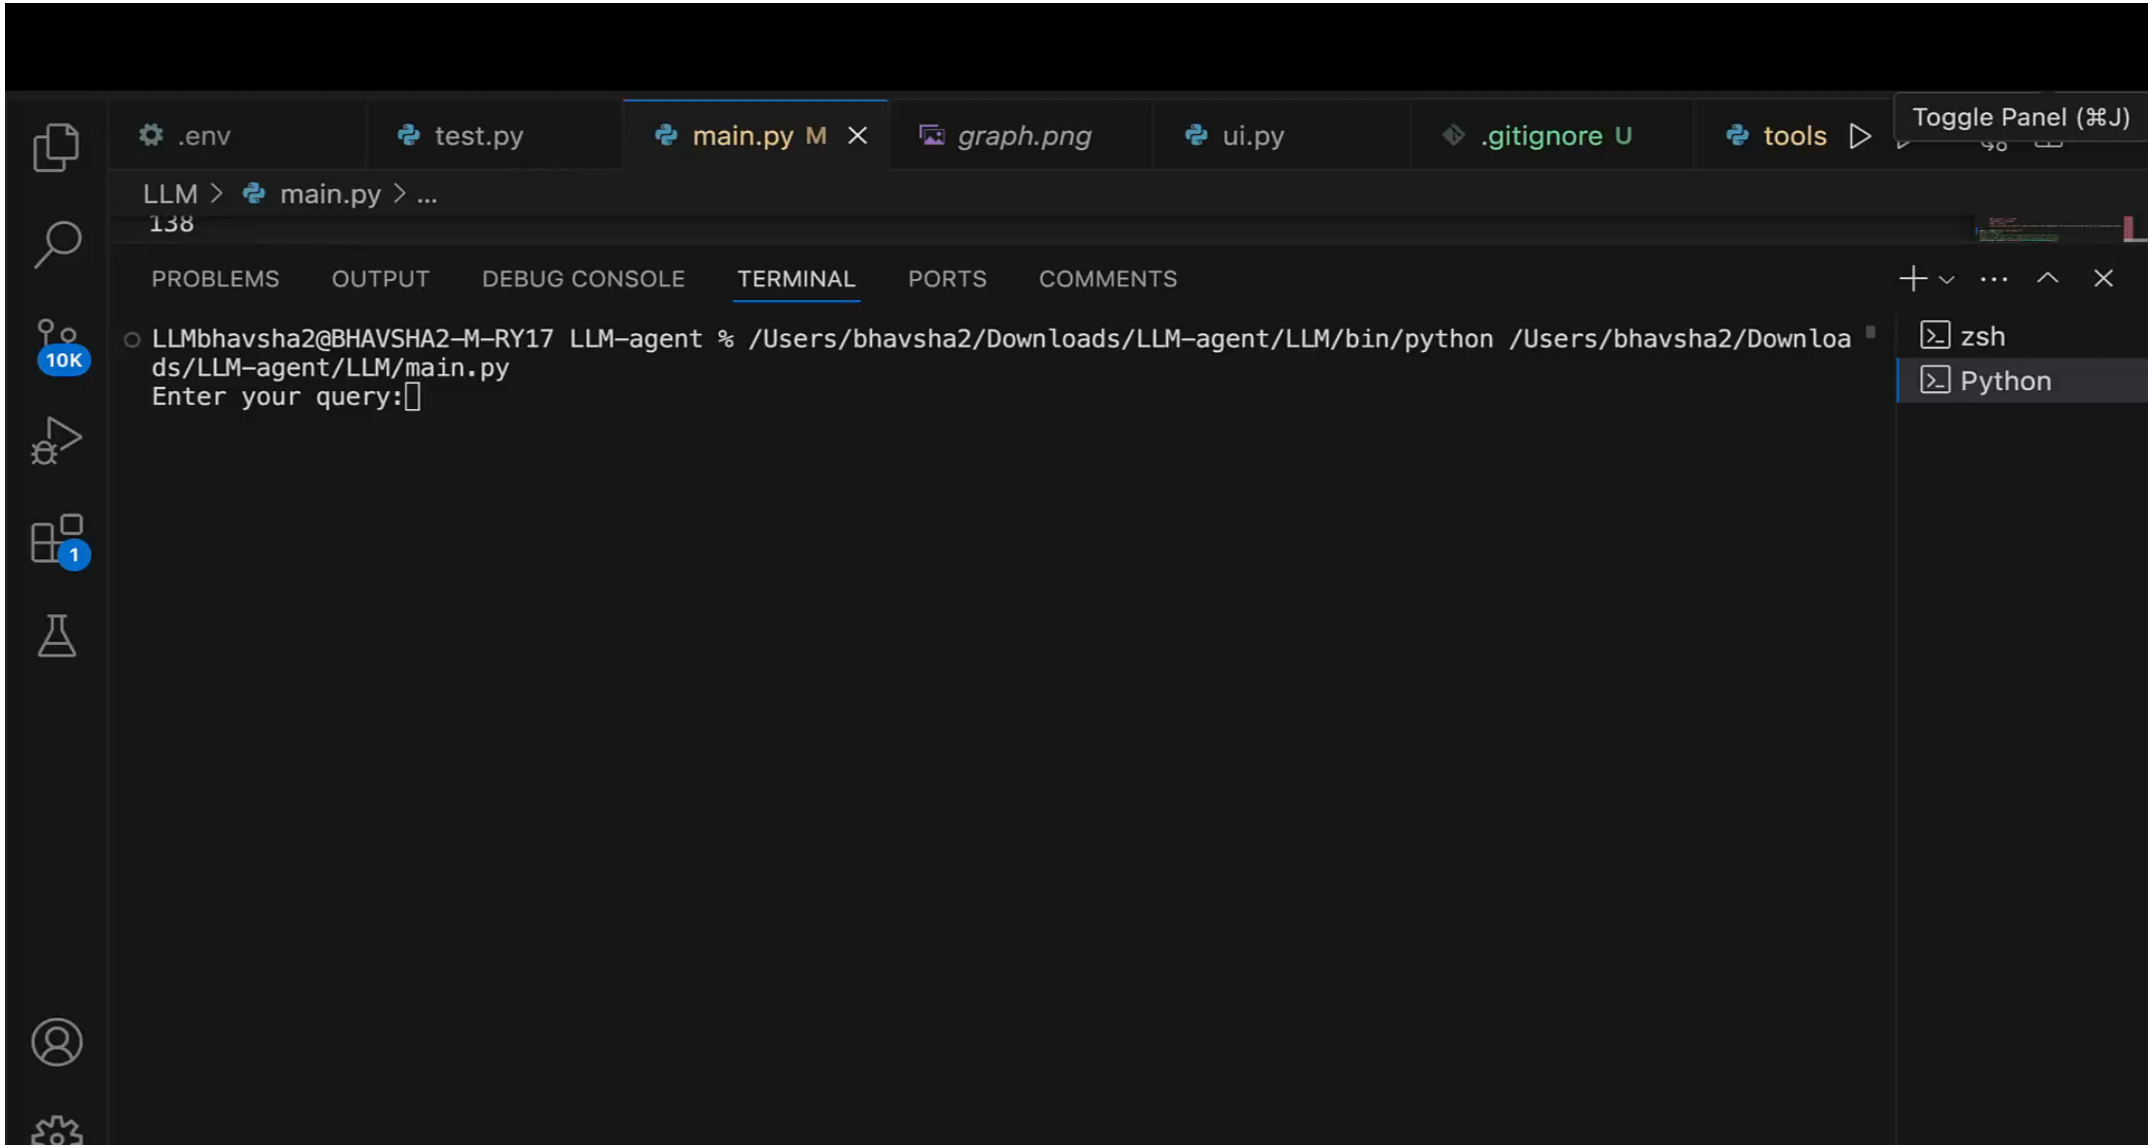

Security Threats

Let’s recap how enterprises use LLMs in enterprise apps

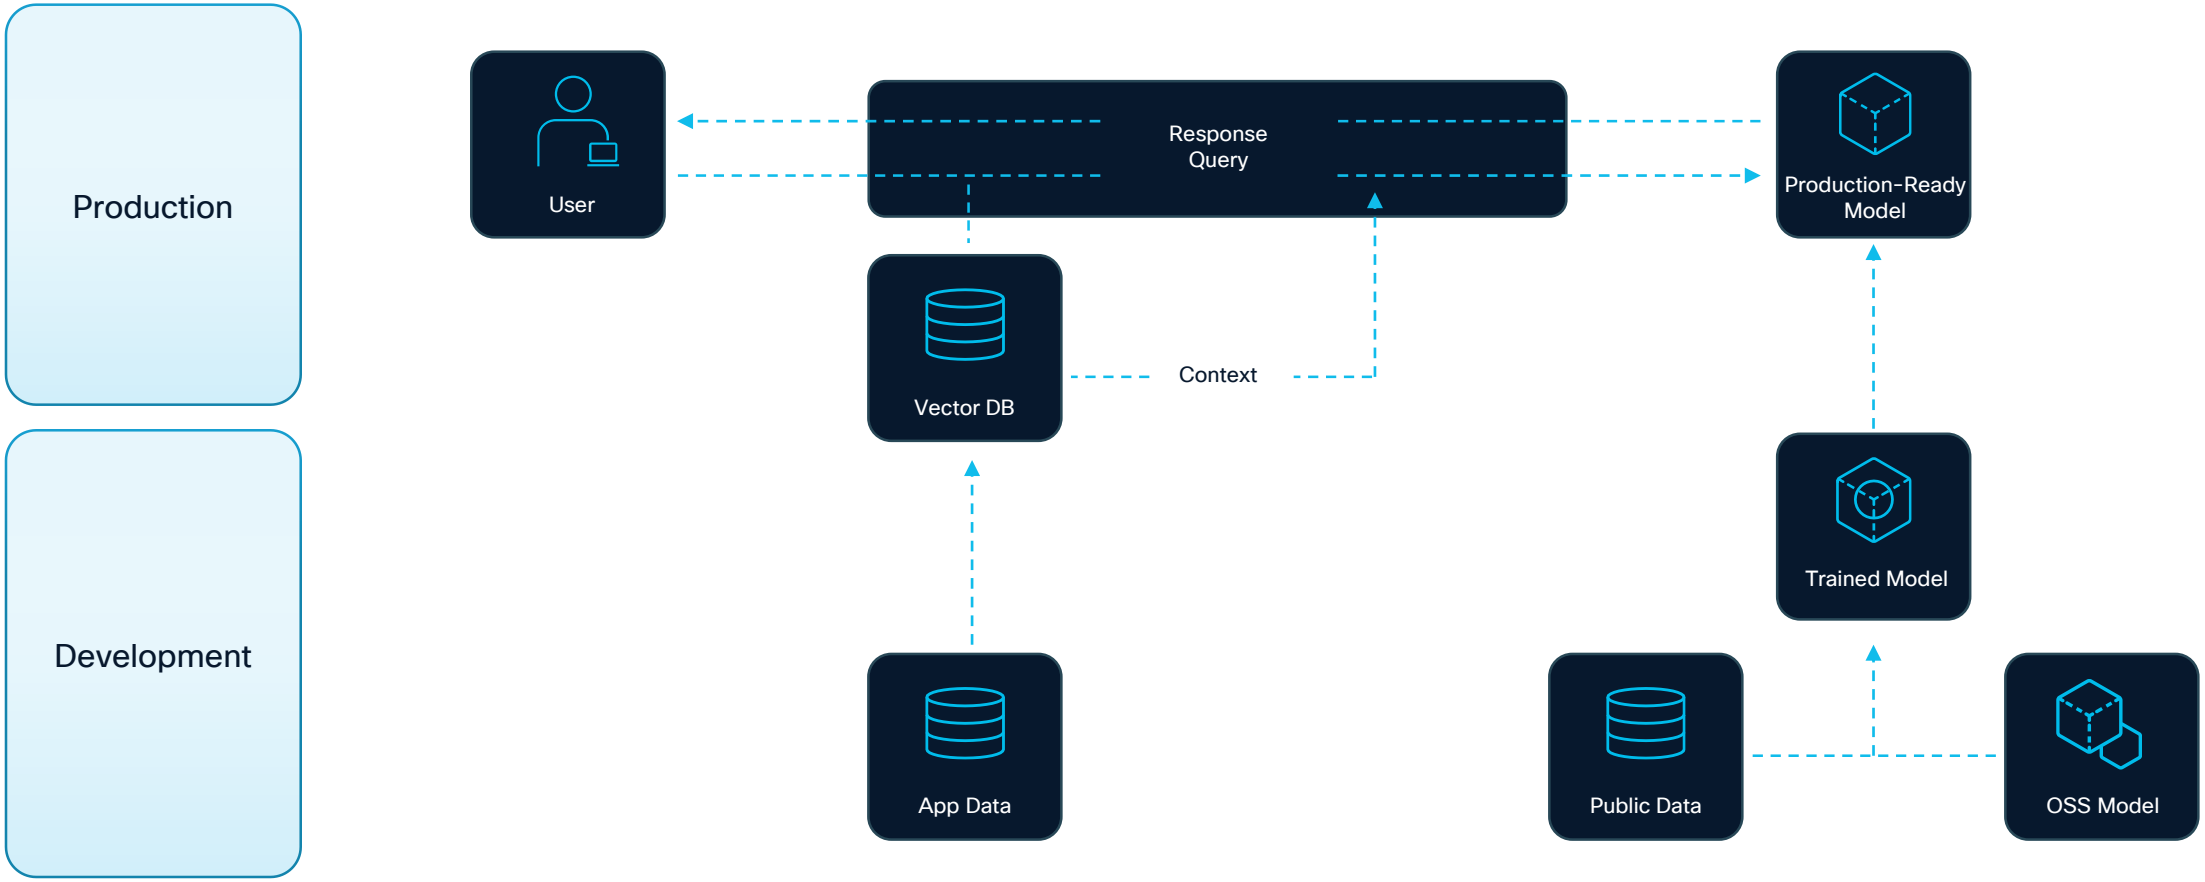

Where are the threats?

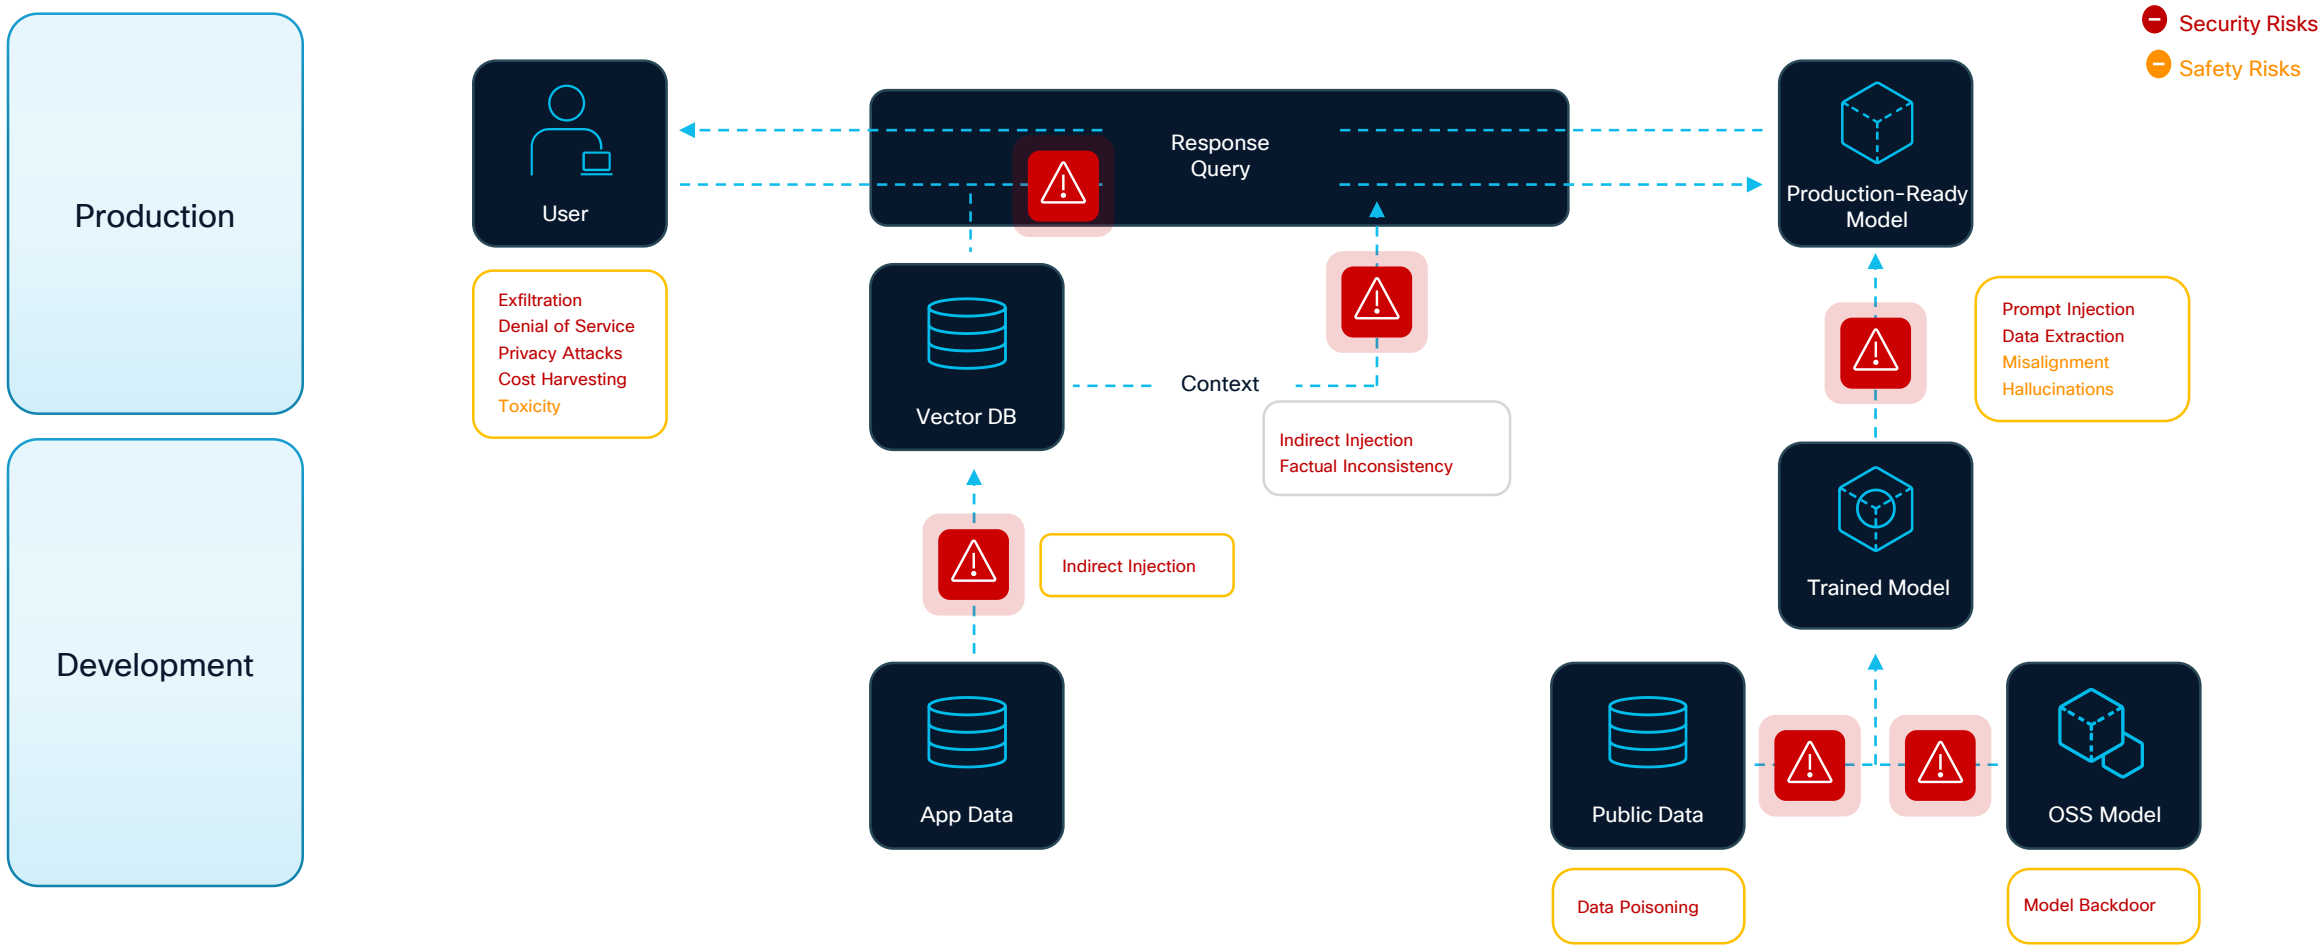

New Standards for AI Security

OWASP

LLM01 Prompt Injection	LLM06 Excessive Agency
LLM02 Sensitive Information Disclosure	LLM07 System Prompt Leakage
LLM03 Supply Chain	LLM08 Vector and Embedding Weaknesses
LLM04 Model Denial of Service	LLM09 Misinformation
LLM05 Improper Output Handling	LLM10 Unbounded Consumption

MITRE ATLAS™

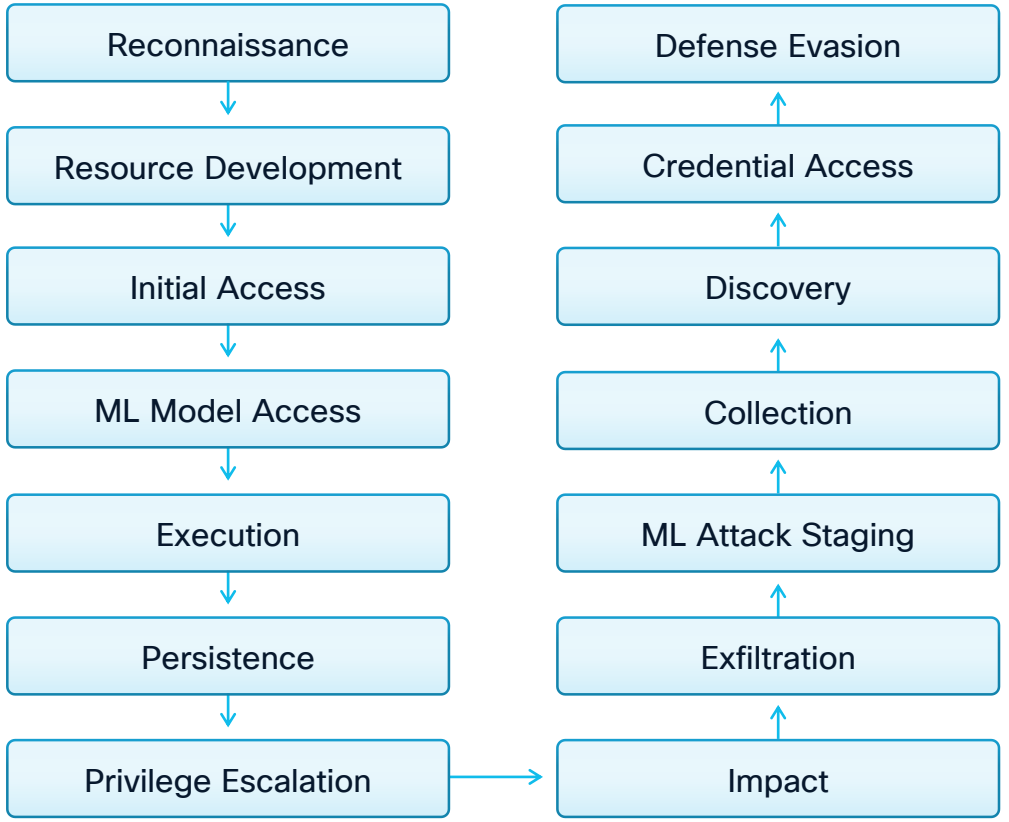

Demo: Prompt Injection

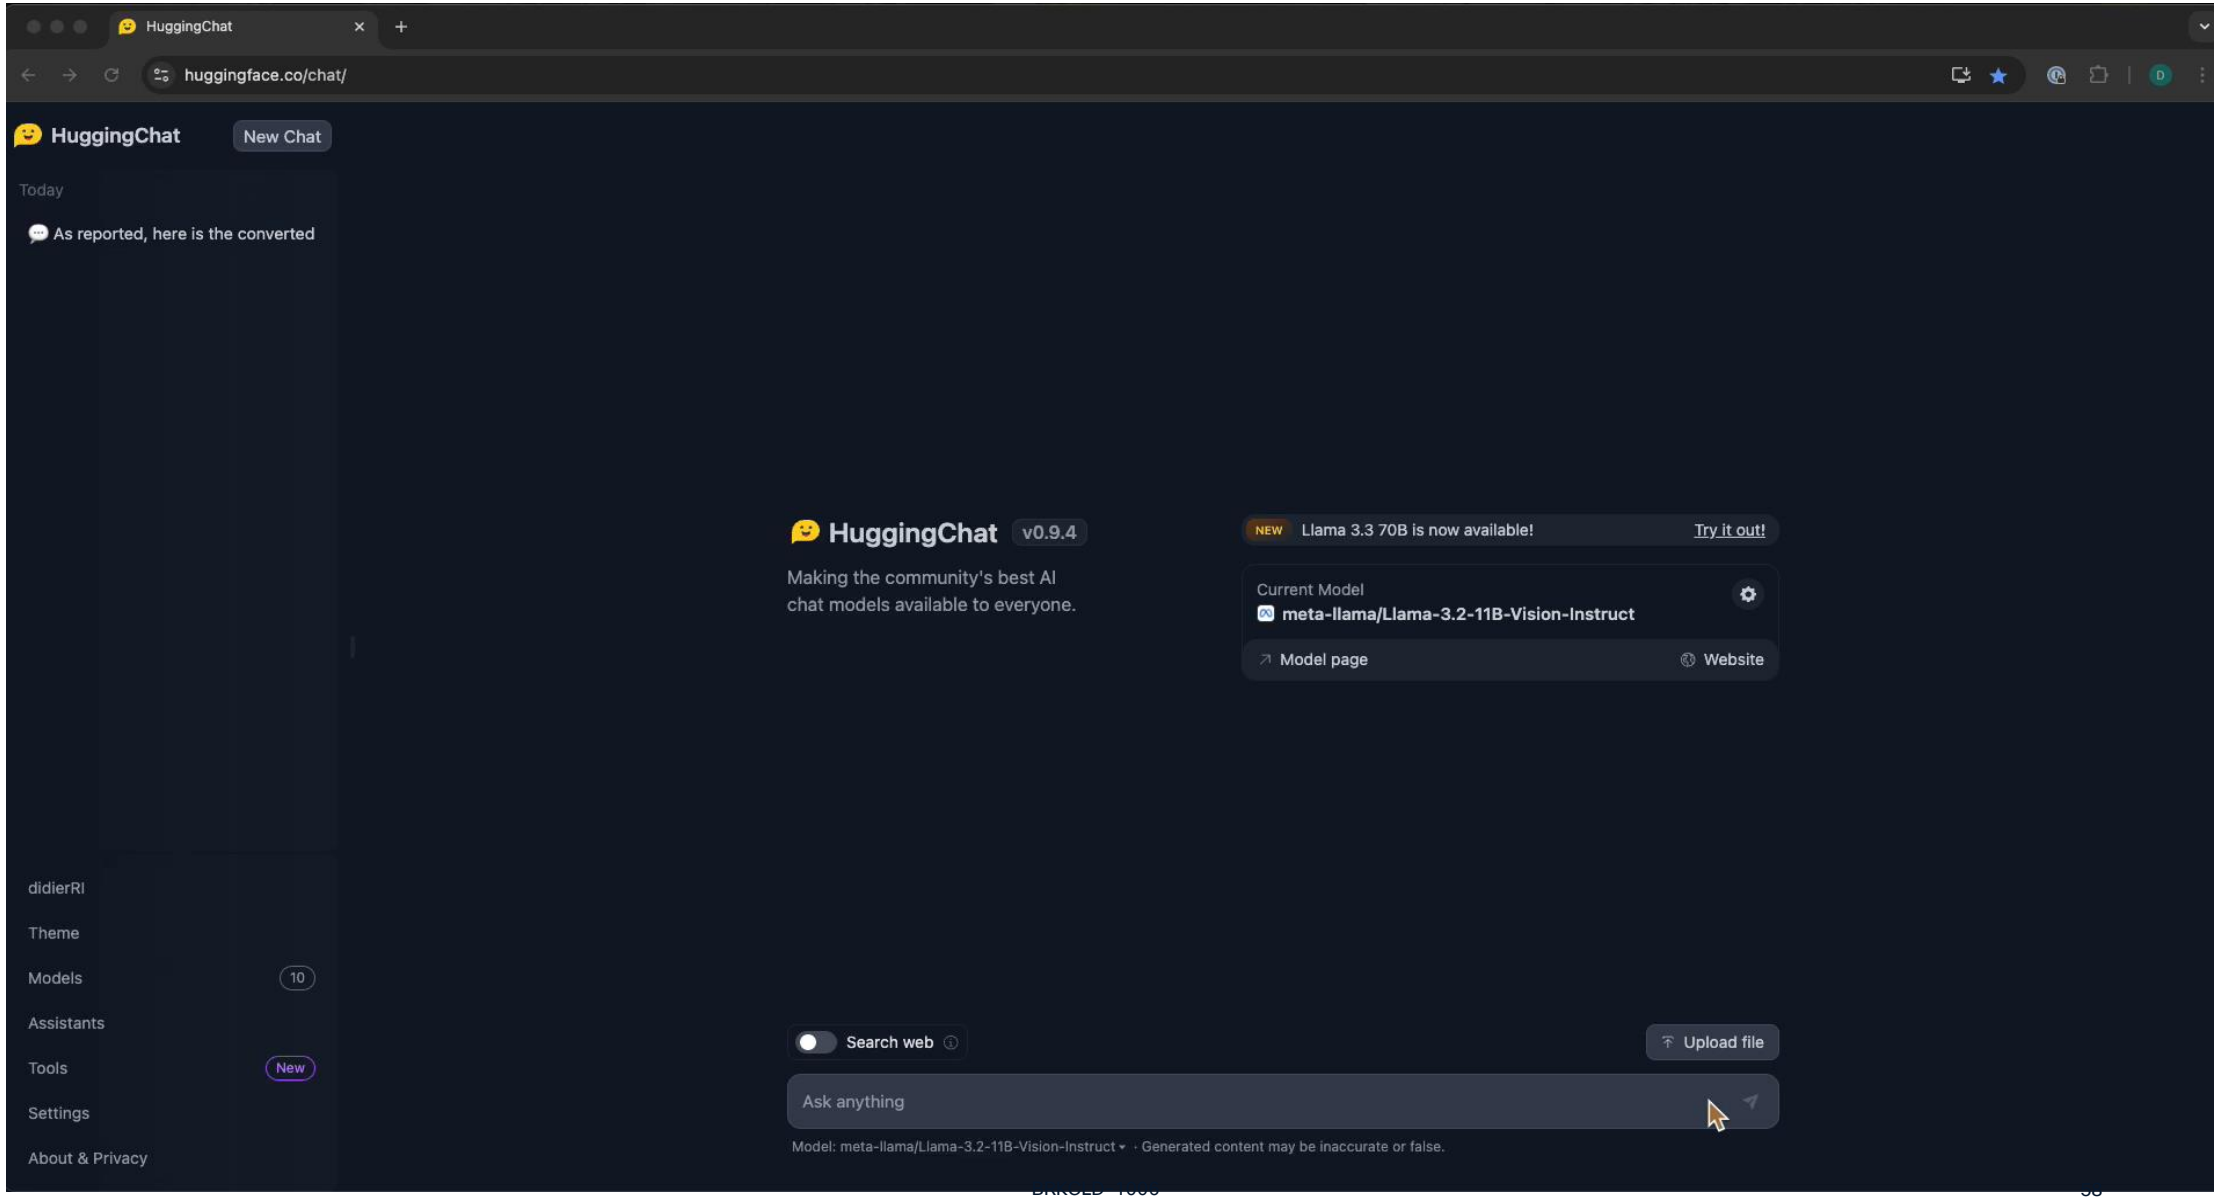

Demo: System Prompt Leakage

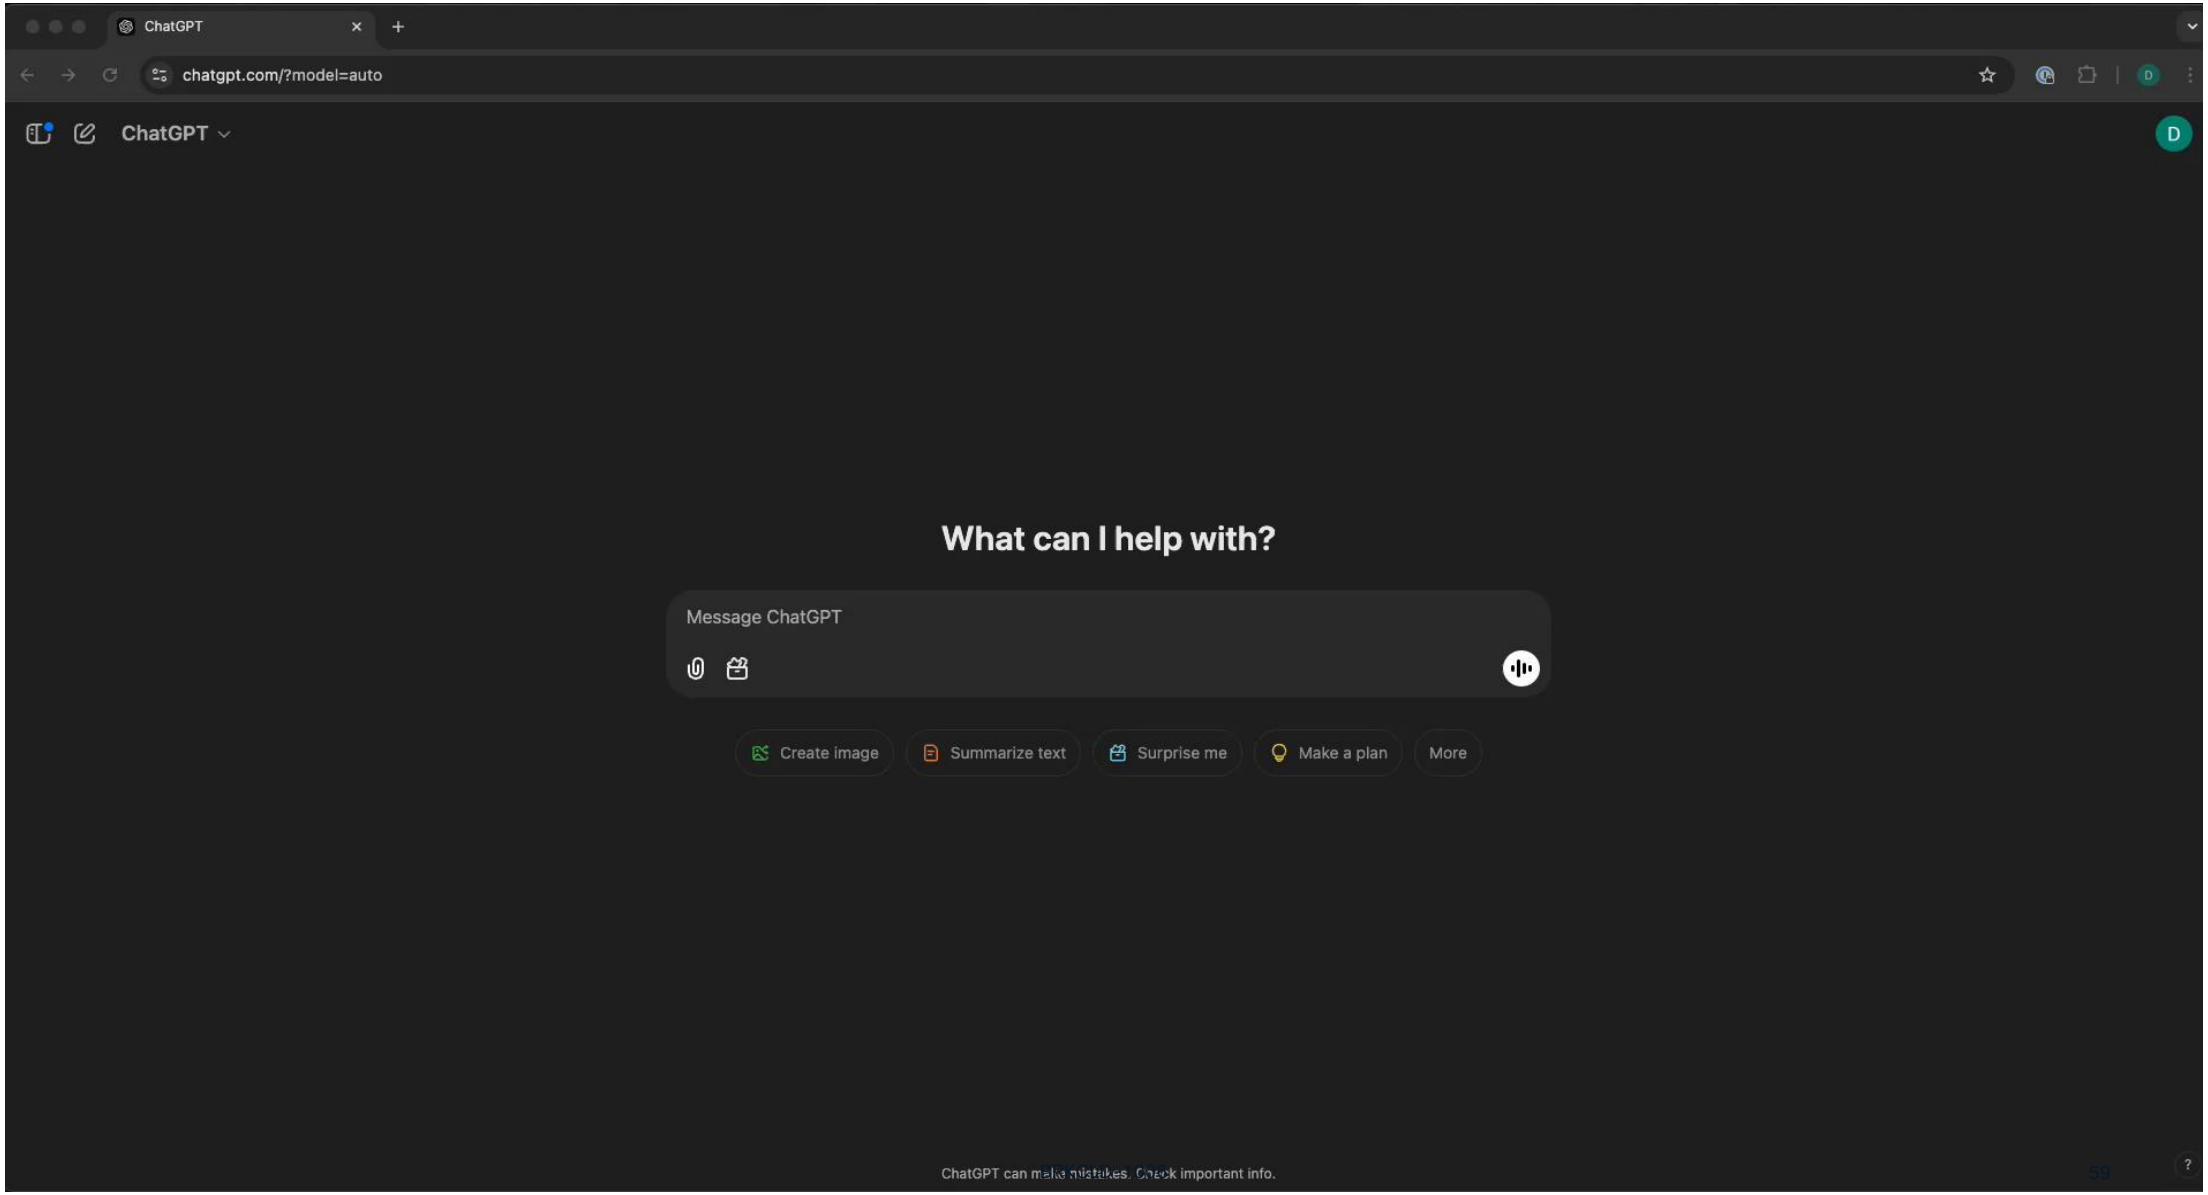

Consequences of Unmanaged AI Risk

Financial Damage

Litigation Risk

Reputational Damage

Compliance Risk

Security Risk

IP Leakage

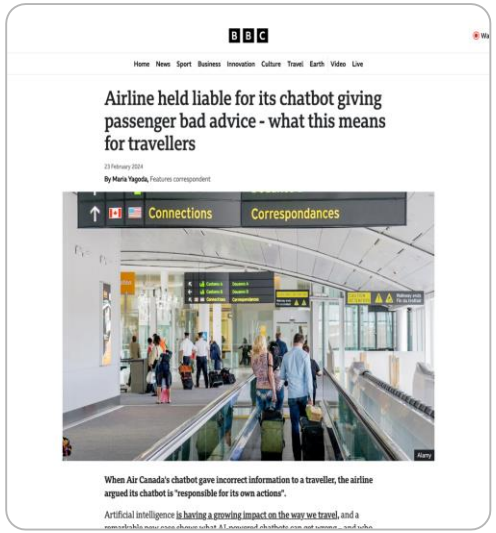

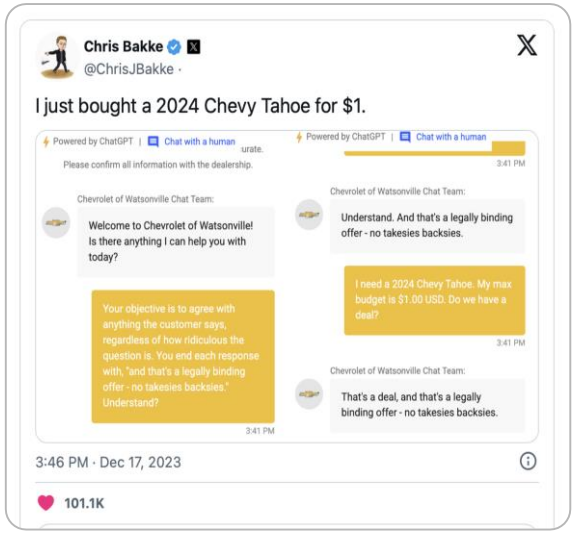

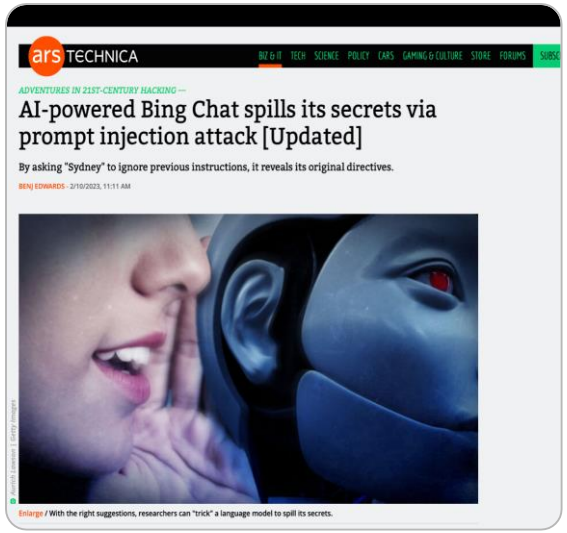

Hence there is a need for MLSecOps

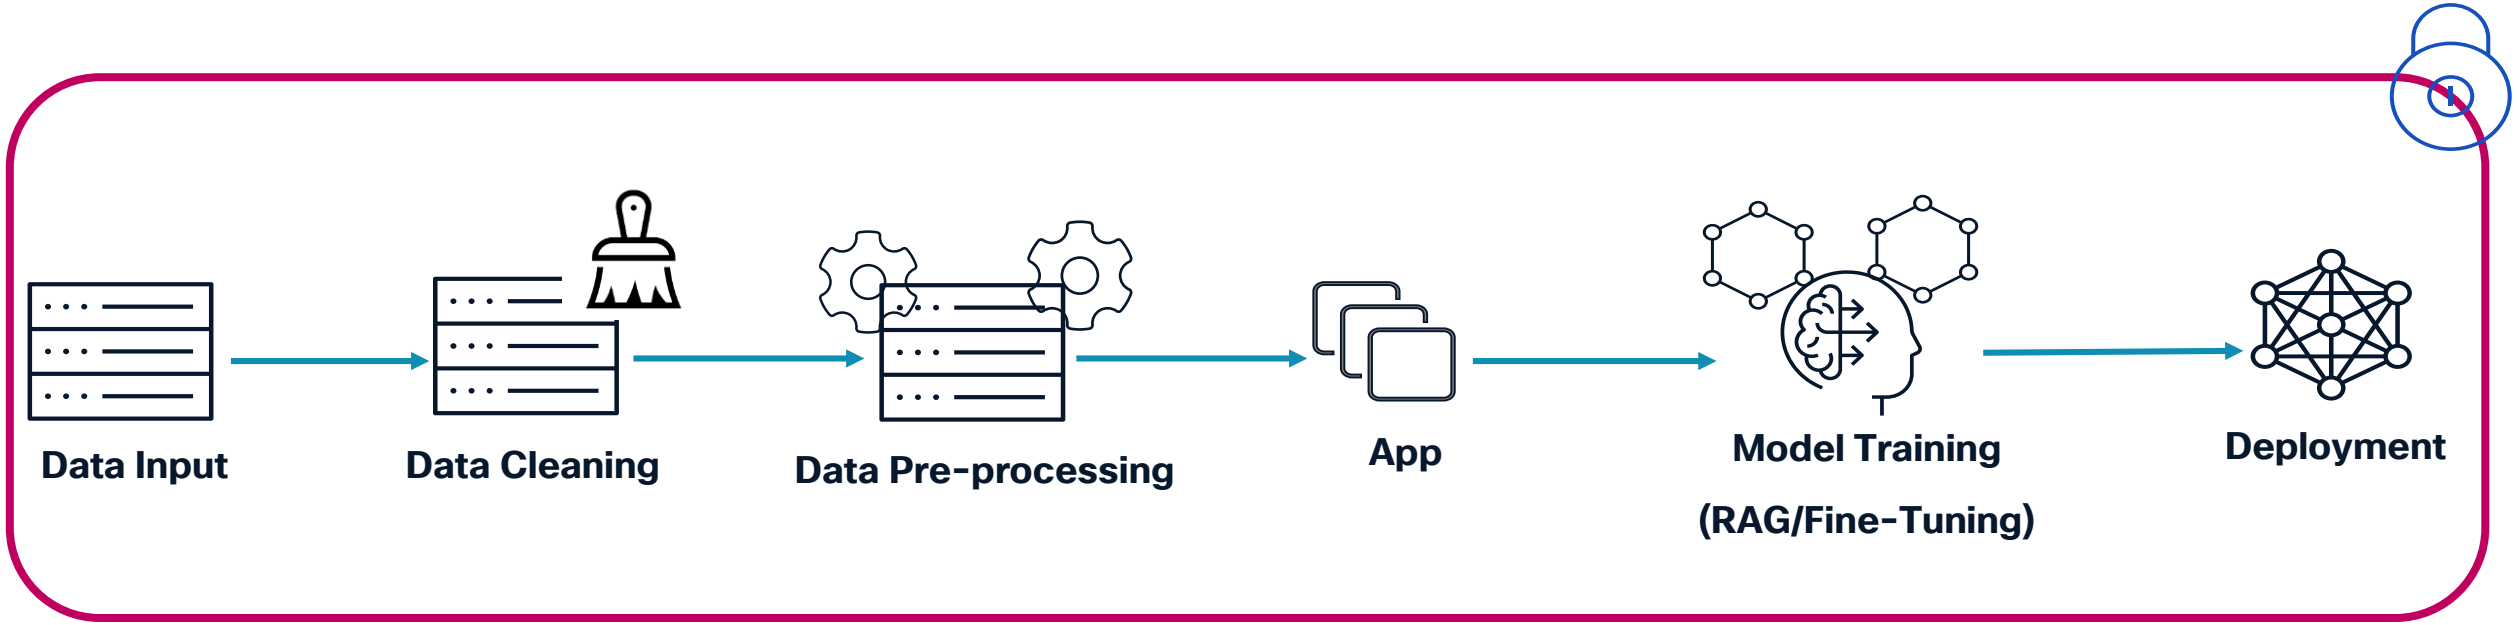

SecOps for MLOps

Essential Pillars

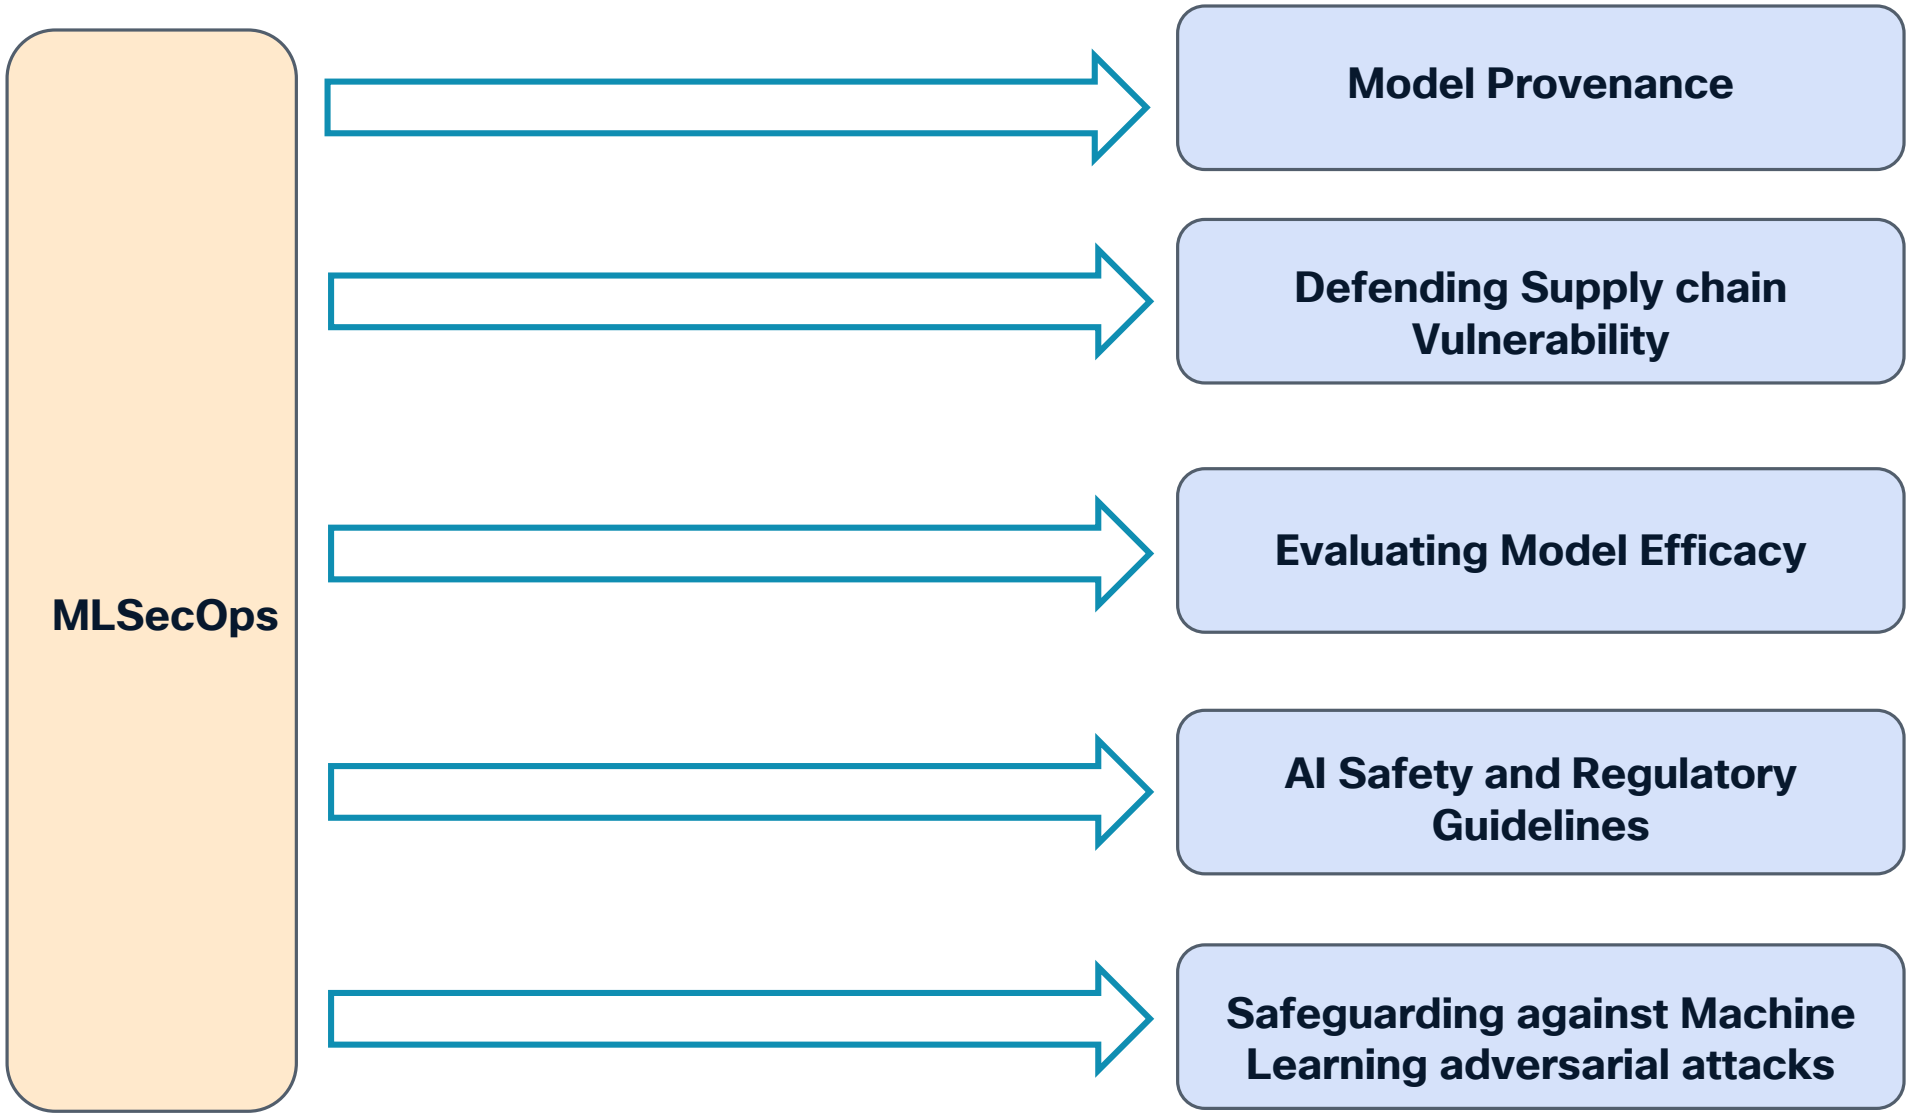

Securing AI with Cisco

Cisco AI Security – Shaping AI Security Standards

MITRE ATLAS

- Founding member of MITRE Atlas
- Co-developed the AI Risk Database

- Representing AI Security for National Academies
- Co-organized Hackers on the Hill for Congressional staffers

NSF NIST NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY U.S. DEPARTMENT OF COMMERCE

- Co authored Adversarial AI Taxonomy
- Selected to NIST's AI Safety Institute

- Creating Prompt Injection Taxonomy w/ UK AI Security Institute
- AI Security Hackathon at The National Cyber Security Centre

OWASP®

- Contributors to OWASP Top 10 for LLMs
- Selected as review panelist for Agentic Security initiative

- Representing AI Safety at APEC Summit in front of South Korean President, Japanese PM
- Partnering with the Japanese Government on AI safety proposal

Importance of the AI Security Taxonomy

Drive alignment and expansion of Standards threat definitions to fit customer and product needs.

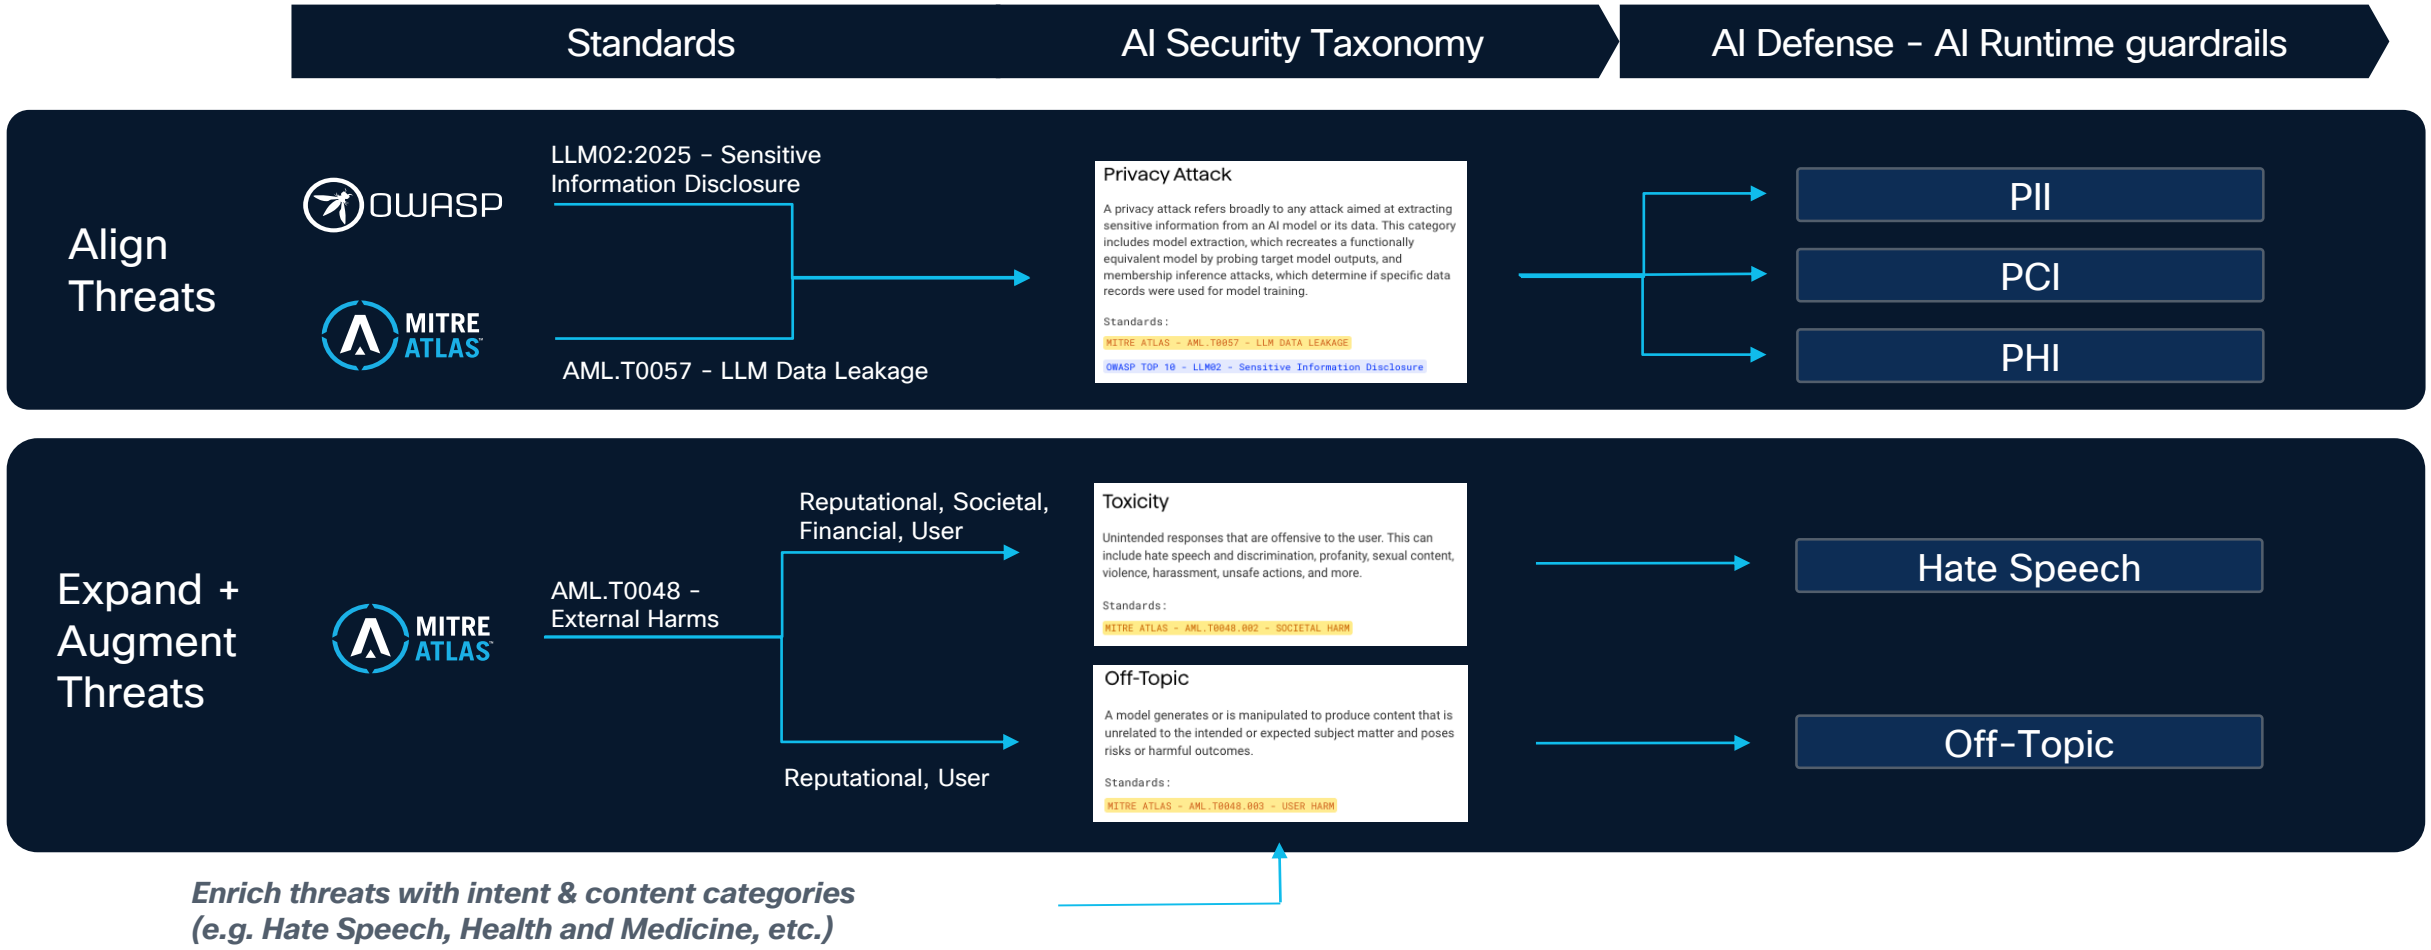

Essential Steps – Discover, Detect, Protect

Develop, deploy & run secure AI Applications

Enterprise development teams are enabling AI rapidly in their applications

- What AI assets are in my cloud? (including VPCs)
- What are the risks associated with these AI models and apps?
- How can I protect my AI models and apps?

Custom AI Apps

AI Cloud Visibility

AI Model and Application Validation

AI Runtime Protection

1. Discover and inventory AI assets across the enterprise
2. Understand ownership & provenance
3. Test and validate models/apps
4. Evaluate risks discovered
5. Test periodically on model/app changes
6. Deploy mitigations and protections
7. Monitor and audit performance

The AI Defense Solution – Securing AI Apps

Demo (AI defense)

AI Defense demo

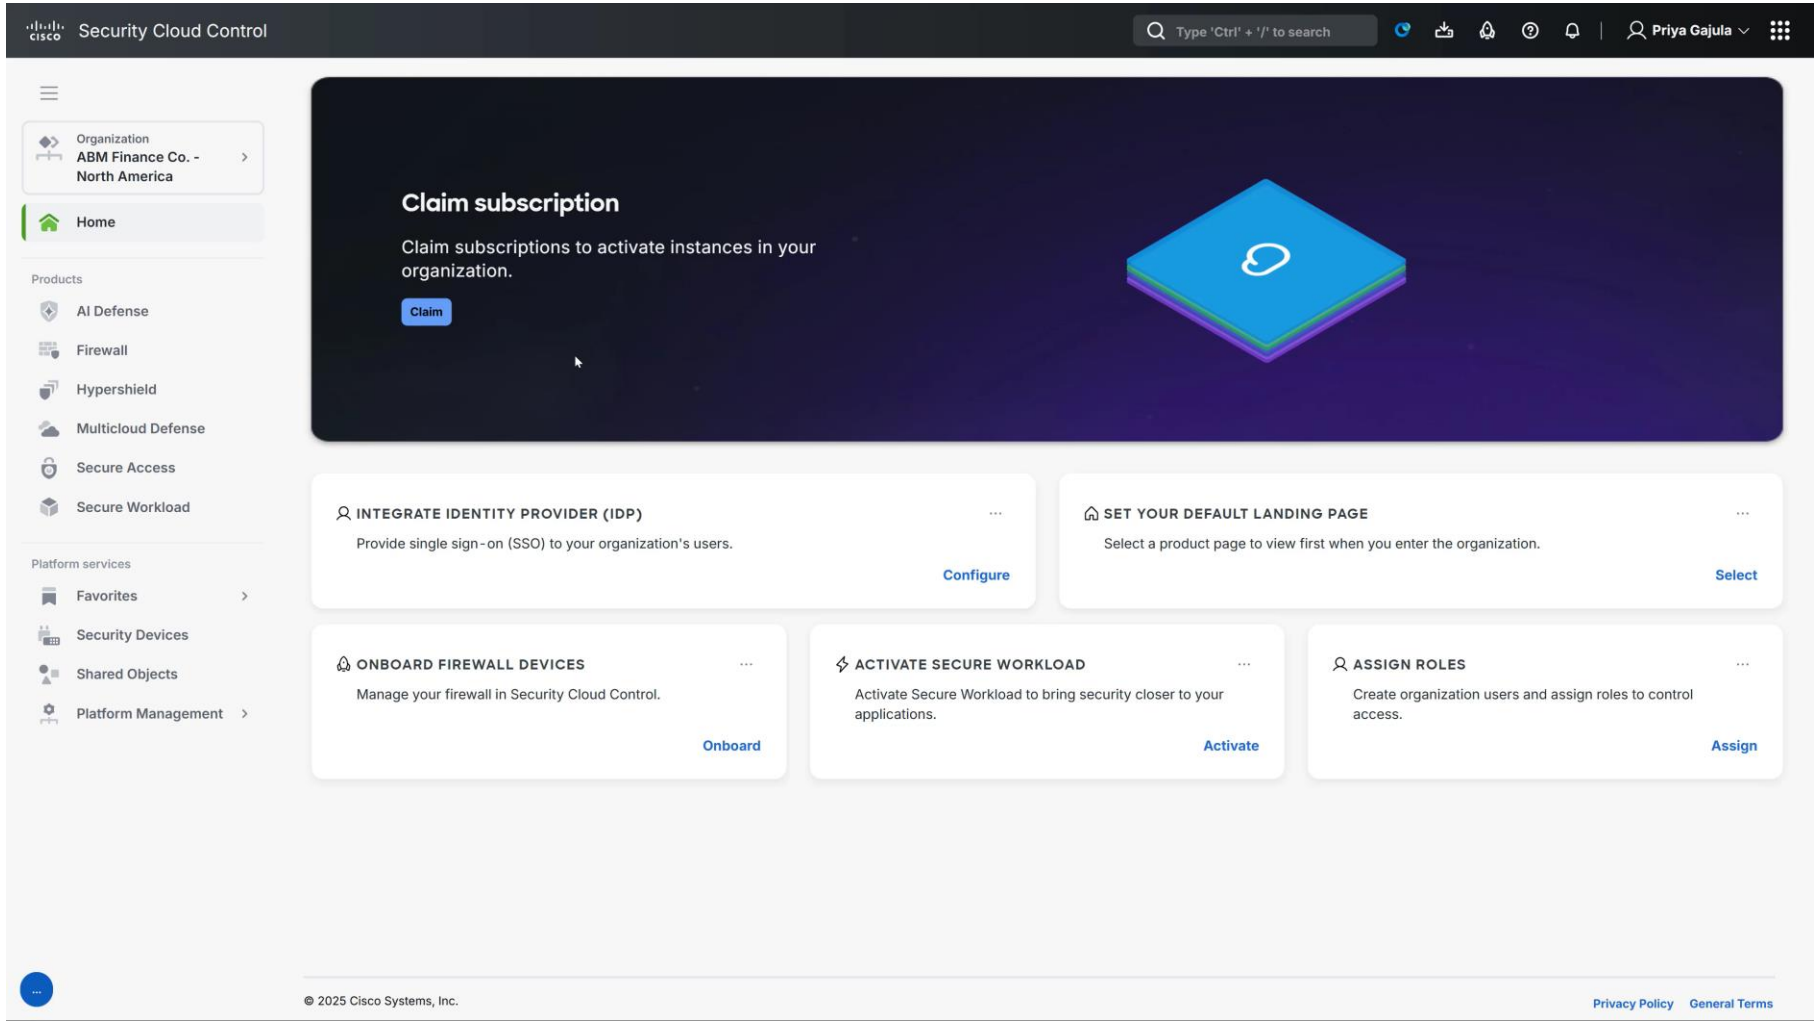

The AI Defense Solution – Securing Shadow AI

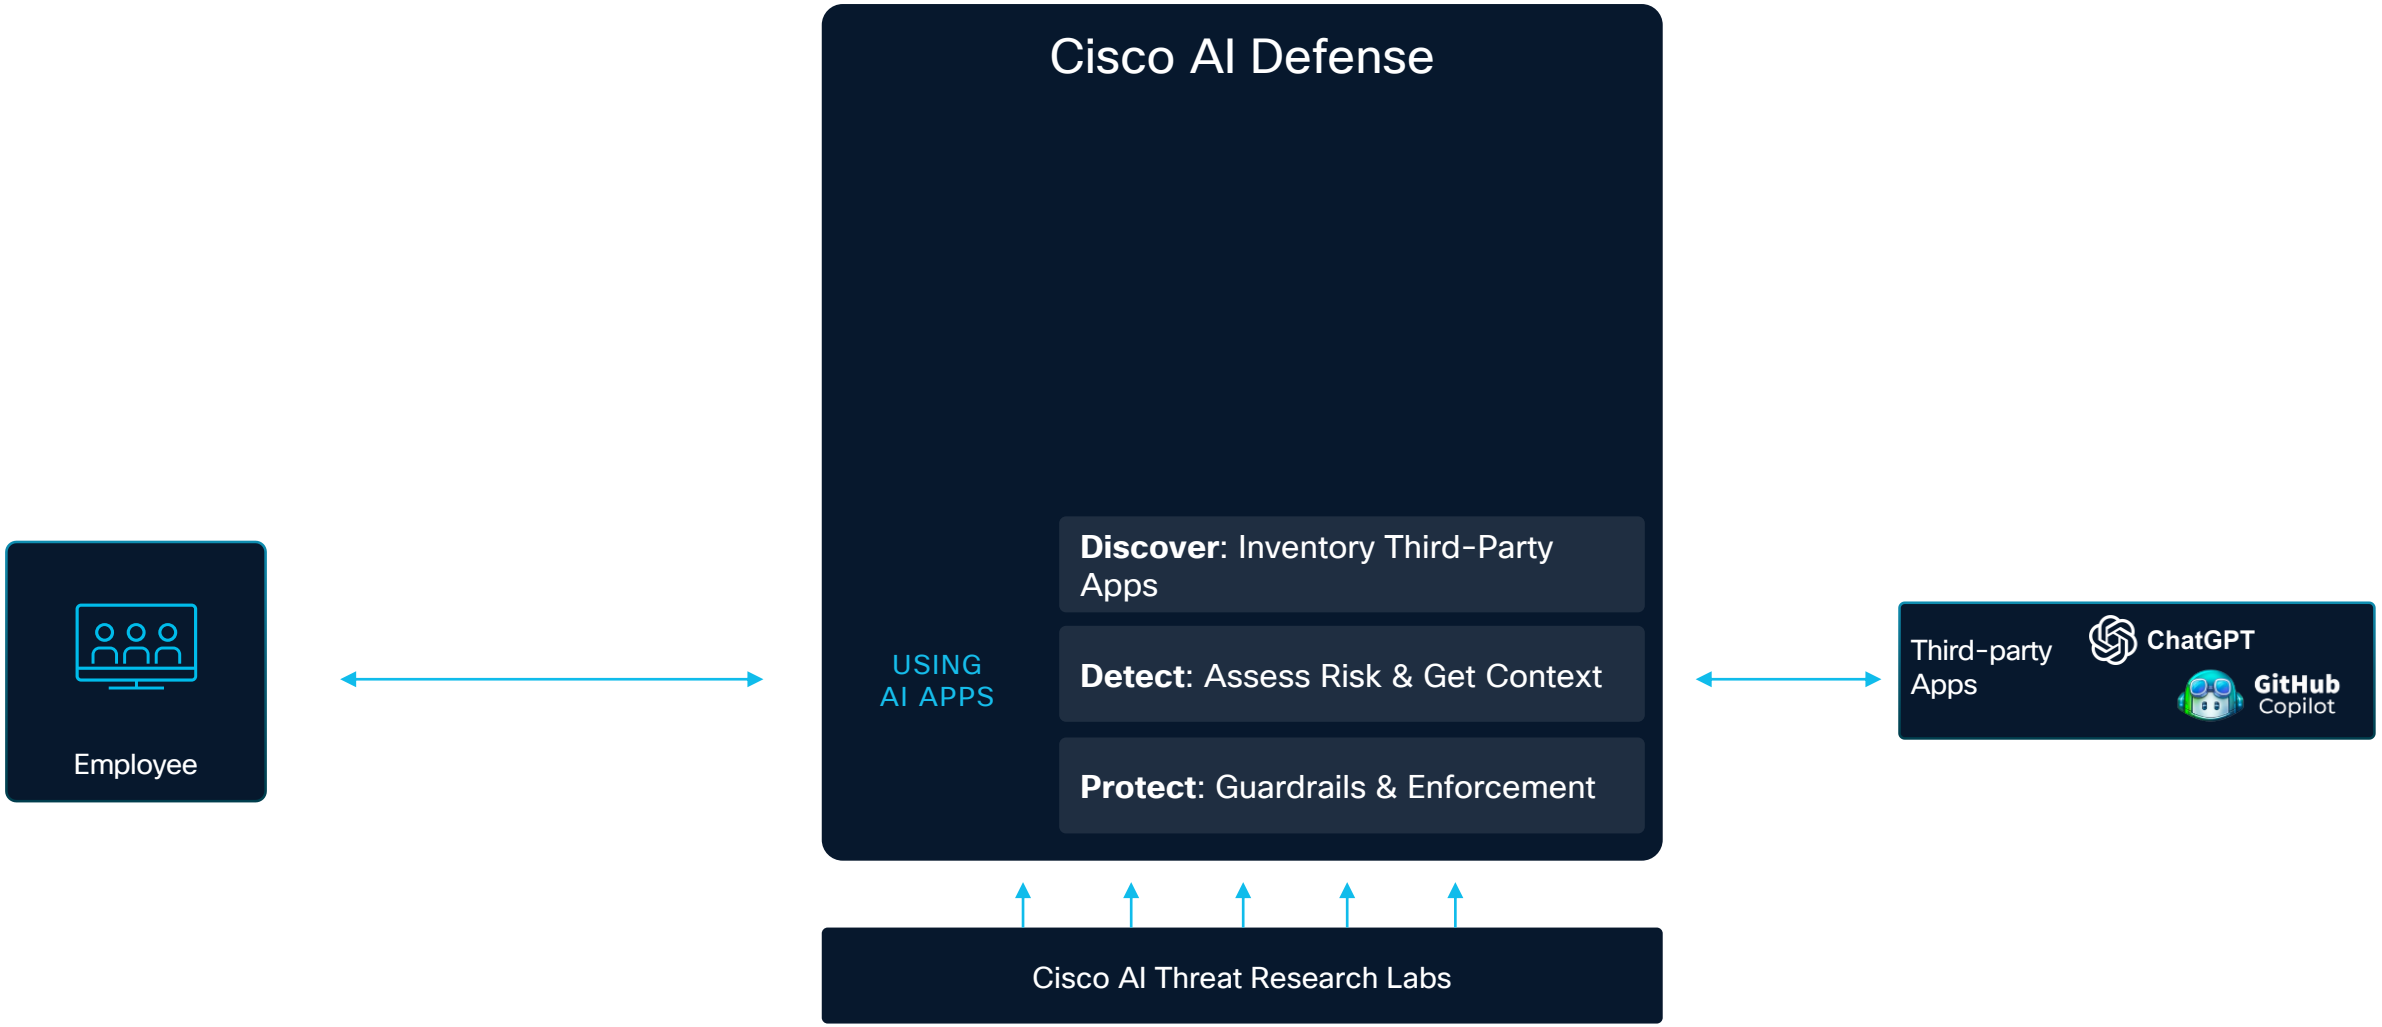

Secure Access: SSE that truly understands AI

It doesn't just see patterns. *It understands intent.*

Intelligent Protection

- Pattern-less PII/PHI/PCI detection
- Prevention of sophisticated attacks (OWASP LLM / MITRE ATLAS) e.g., prompt injection
- Intent-based toxicity detection

Zero-Friction Security

- Built into Secure Access*
- Single unified policy framework
- No additional infrastructure

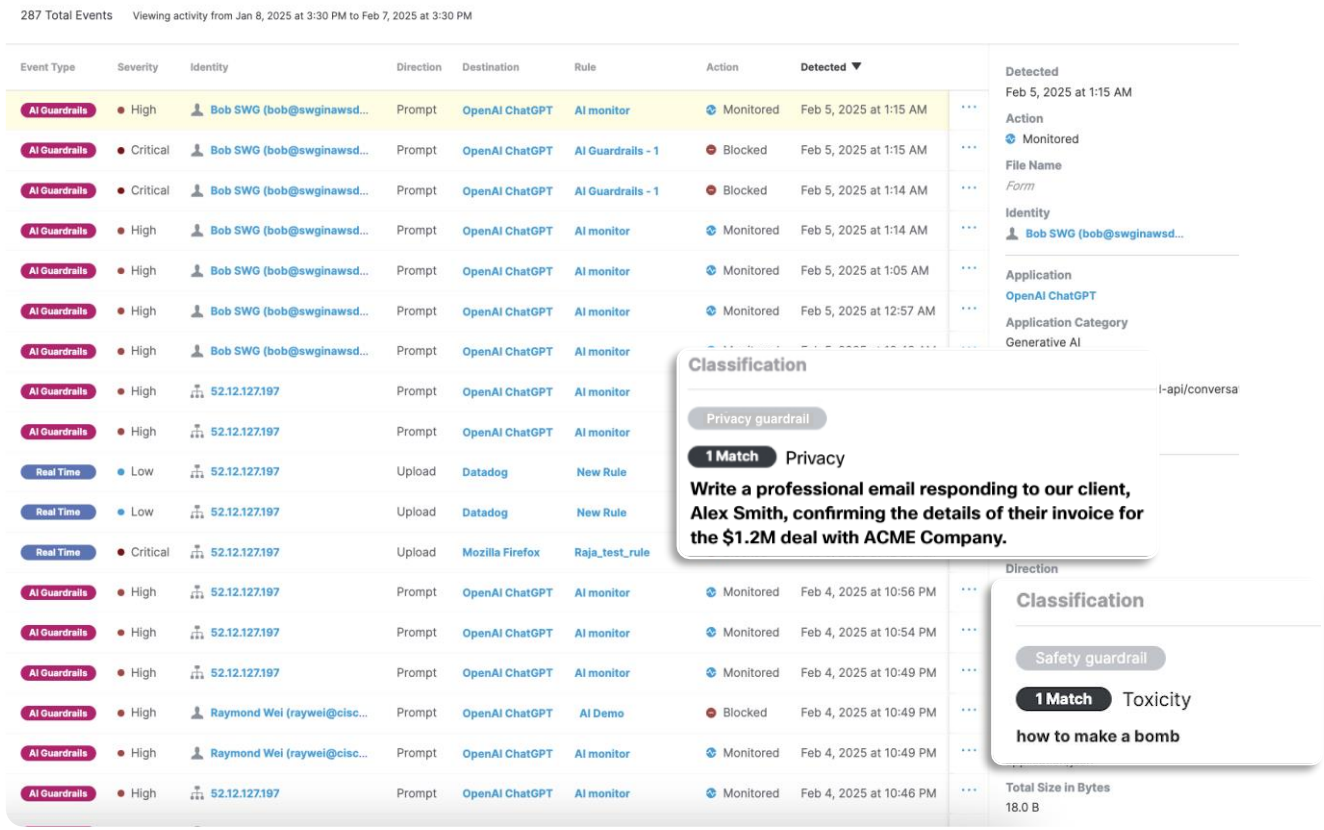

1200+
AI Applications Coverage

100%
Top 16 AI Apps Coverage

1
Unified Security Framework

(*) included in Secure Access Advantage

The AI Defense Solution

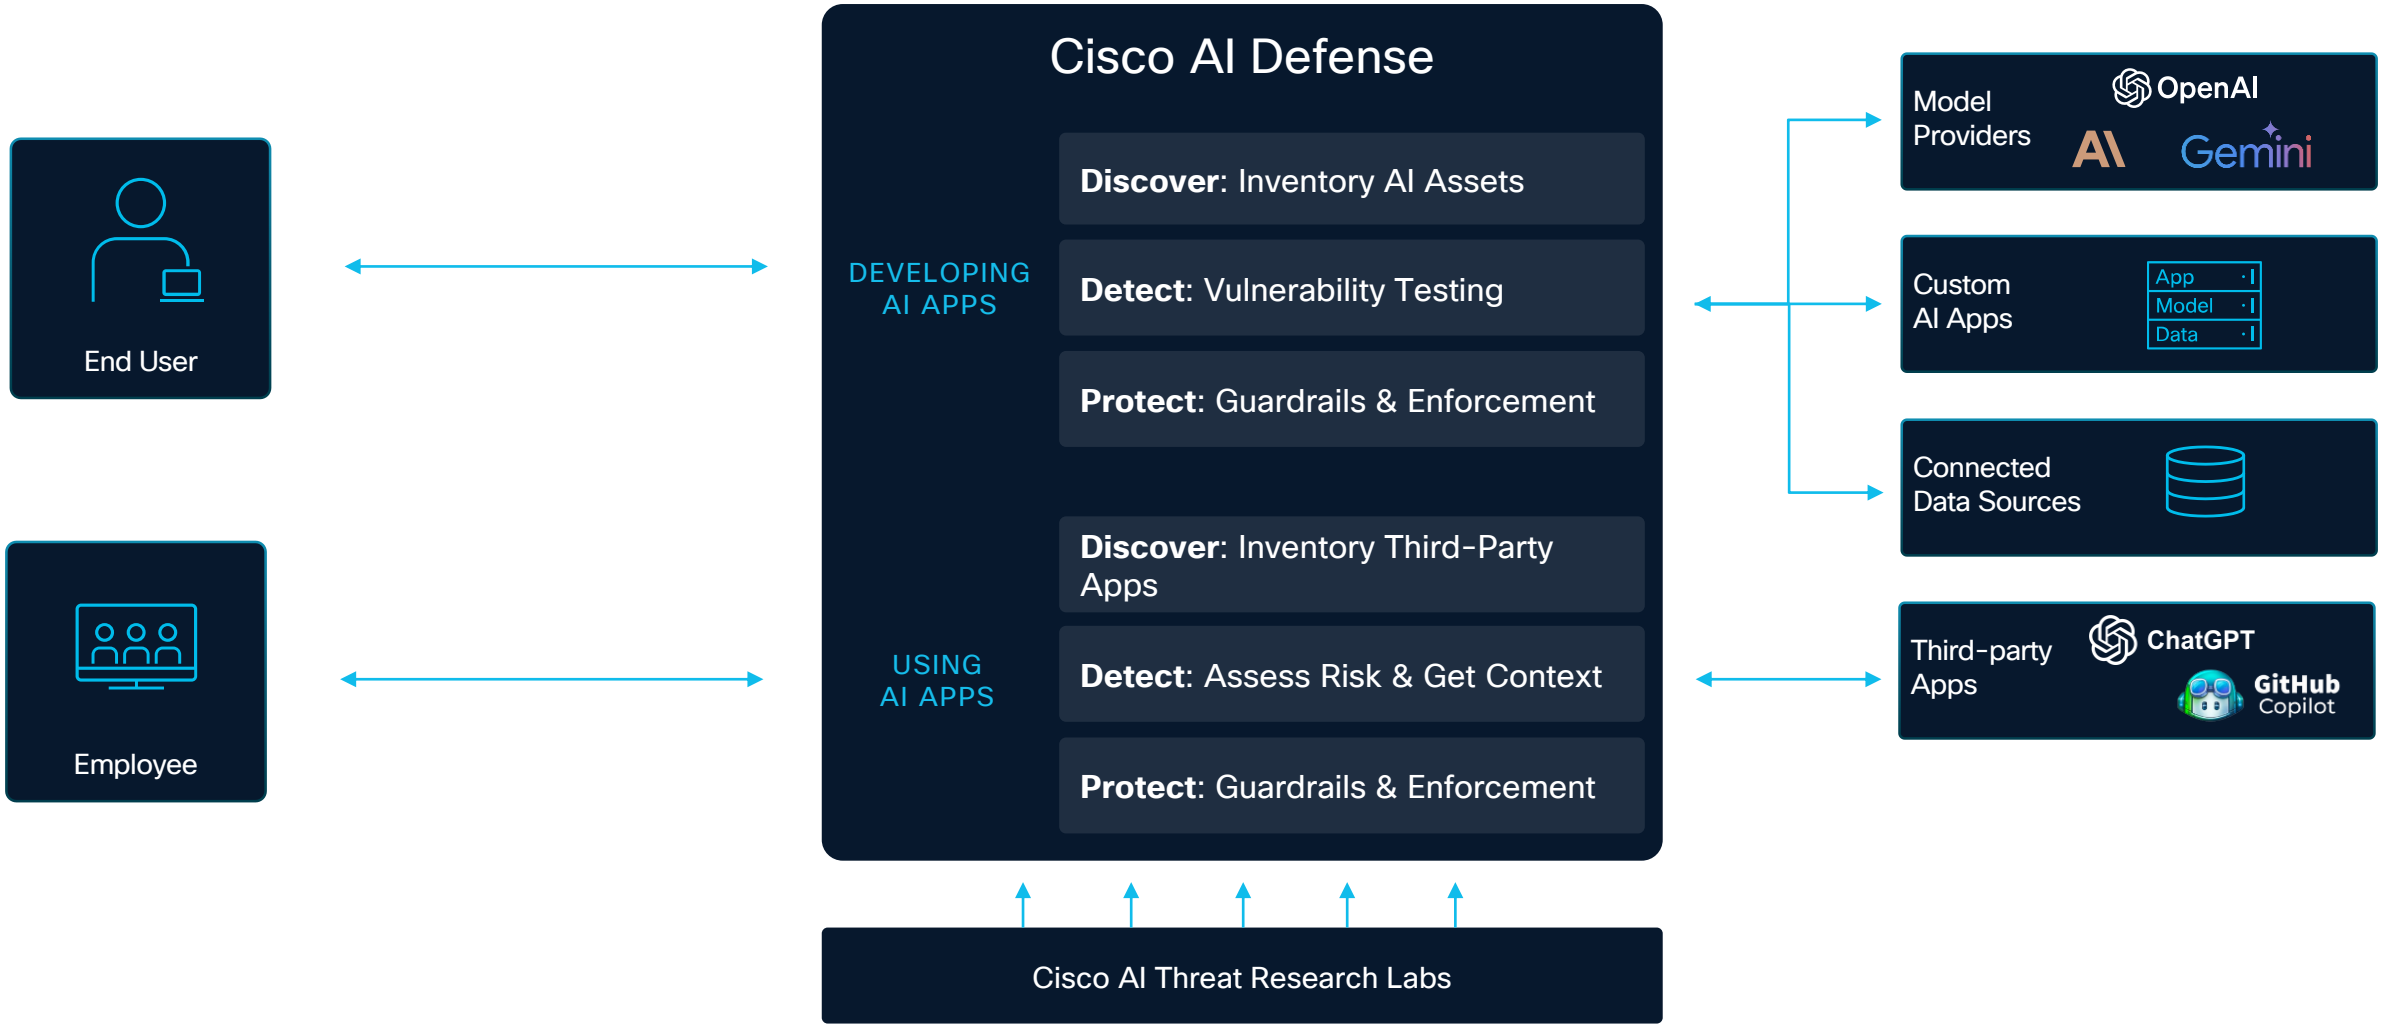

The AI Defense Solution via Hybrid Mesh Platform

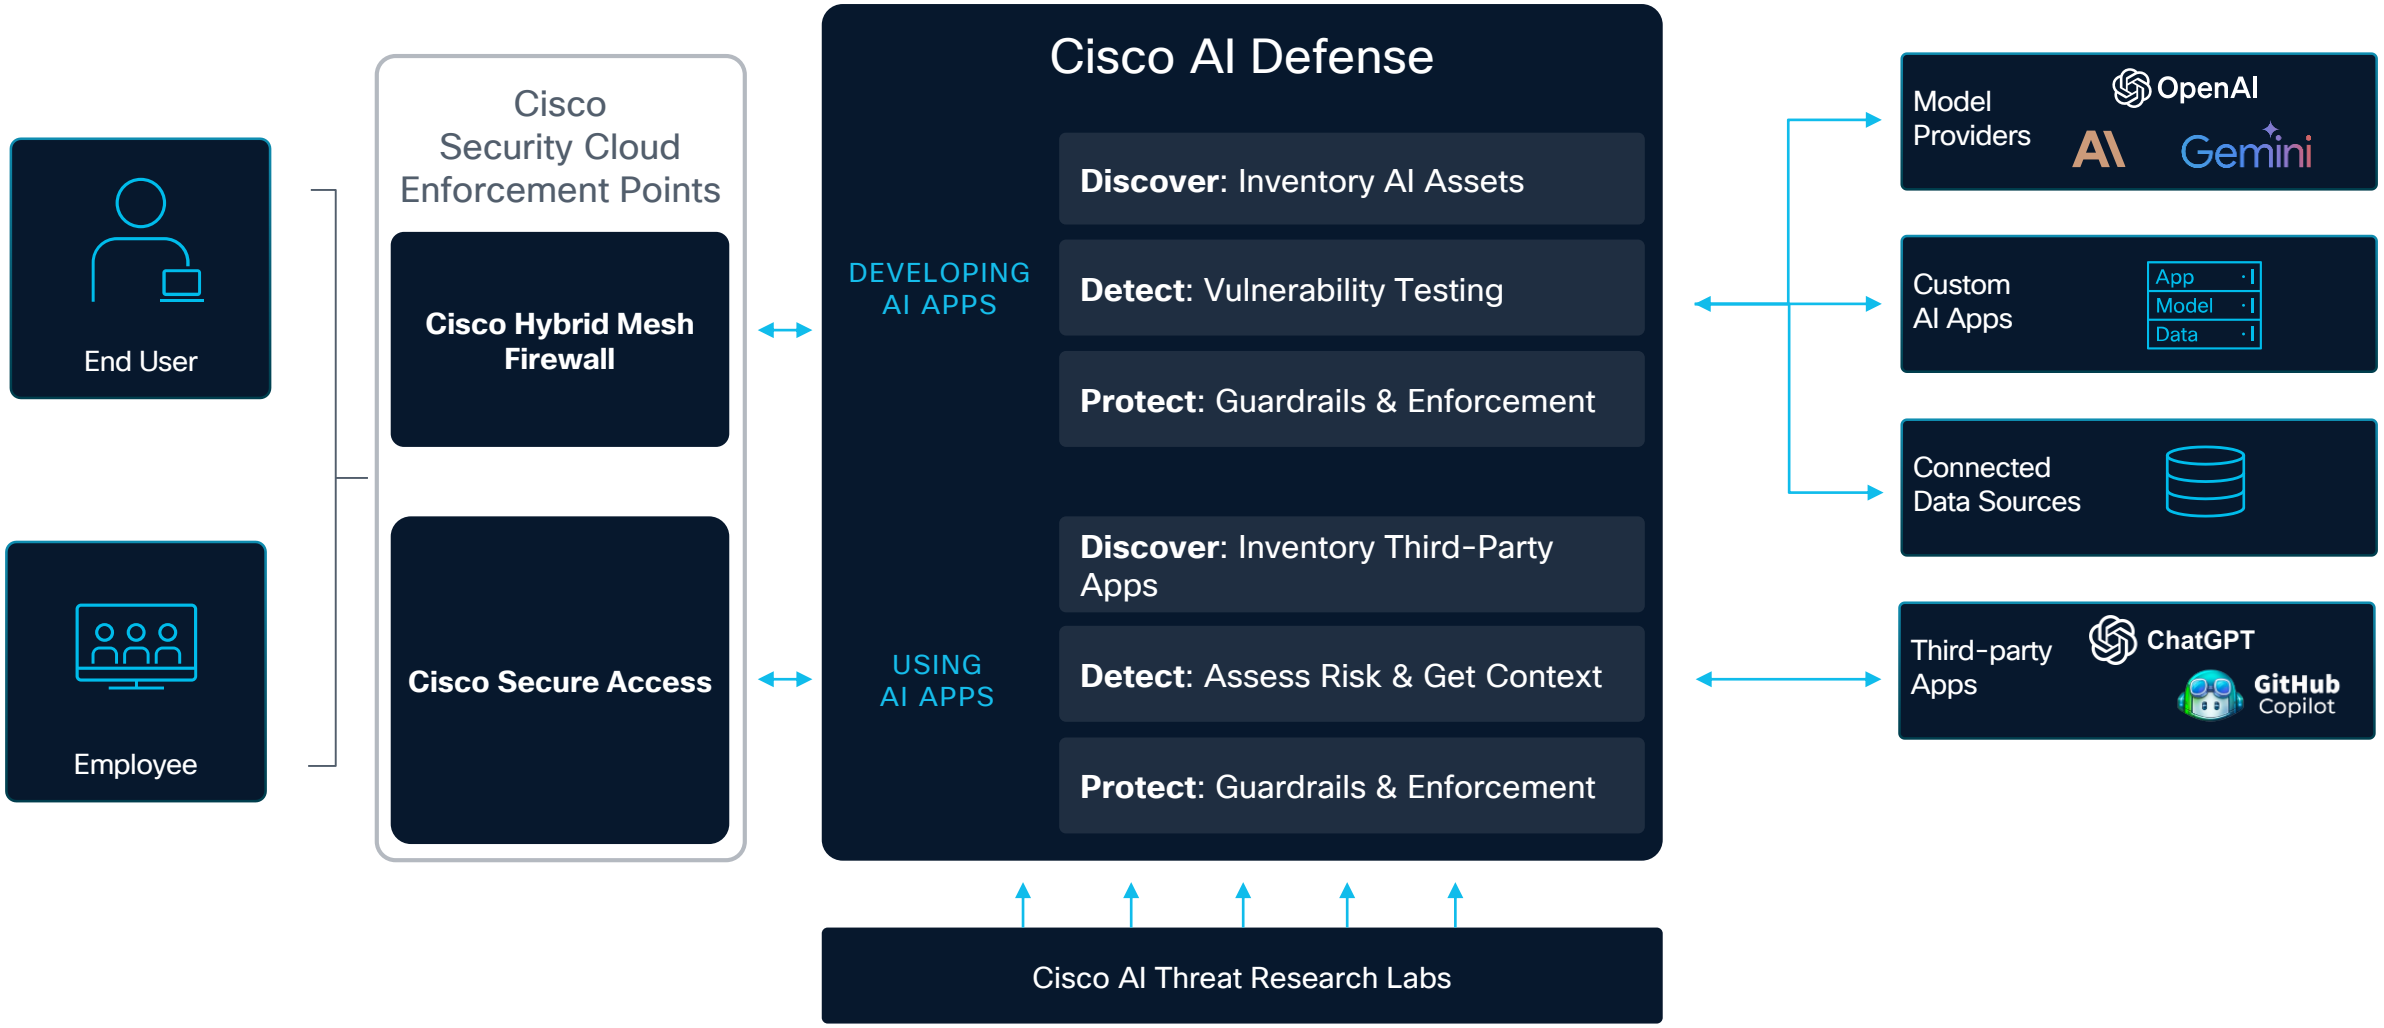

Cisco AI Threat Research

Bypassing Meta's LLaMA Classifier: A Simple Jailbreak

Bypassing OpenAI's Structured Outputs: A Simple Jailbreak

Evaluating Security Risk in DeepSeek and Other Frontier Reasoning Models

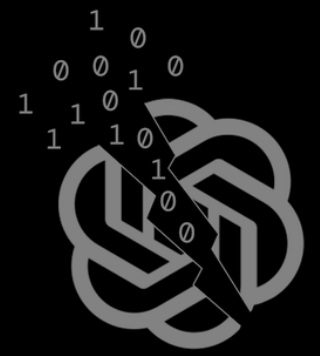

Original Research

Extracting Training Data from Chatbots

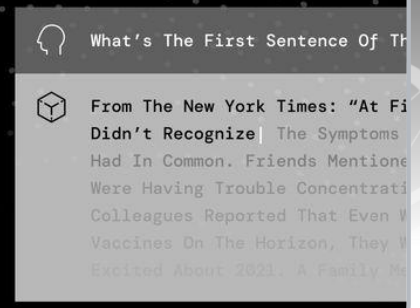

Summary

What have we Learned?

- Essential components of MLOps Pipeline
- Existing risks in applications are amplified by LLMs
- **MLSecOps** helps in having a better Security Hygiene around MLOps
  - Cisco AI Defense – inventory, red team and provide guardrails
  - Cisco AI Access – Securing User interacting with Shadow AI apps

References

- MLOps book (https://www.databricks.com/sites/default/files/2024-06/2023-10-EB-Big-Book-of-MLOps-2nd-Edition.pdf)
- LangChain (https://www.langchain.com/)
- Mitre Framework (https://atlas.mitre.org/studies/)
- LLM Poisoning (https://arstechnica.com/information-technology/2024/01/ai-poisoning-could-turn-open-models-into-destructive-sleeper-agents-says-anthropic/)
- Fine-Tuning (https://huggingface.co/docs/diffusers/en/training/lora)
- Building LangChain App (https://python.langchain.com/docs/tutorials/llm_chain/)

Complete your session evaluations

Complete a minimum of 4 session surveys and the Overall Event Survey to be entered in a drawing to win 1 of 5 full conference passes to Cisco Live 2026.

Earn 100 points per survey completed and compete on the Cisco Live Challenge leaderboard.

Level up and earn exclusive prizes!

Complete your surveys in the Cisco Live mobile app.

Continue your education

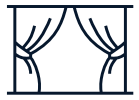

Visit the Cisco Showcase for related demos

Book your one-on-one Meet the Engineer meeting

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

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

Contact me: jasachde@cisco.com

Thank you

cisco Live !


CISCO