What are Agentic AI Threats?

A Cloud Security Perspective

Felix Brehmer
Enterprise Account Executive

Cedric Feist
Senior Solutions Engineer

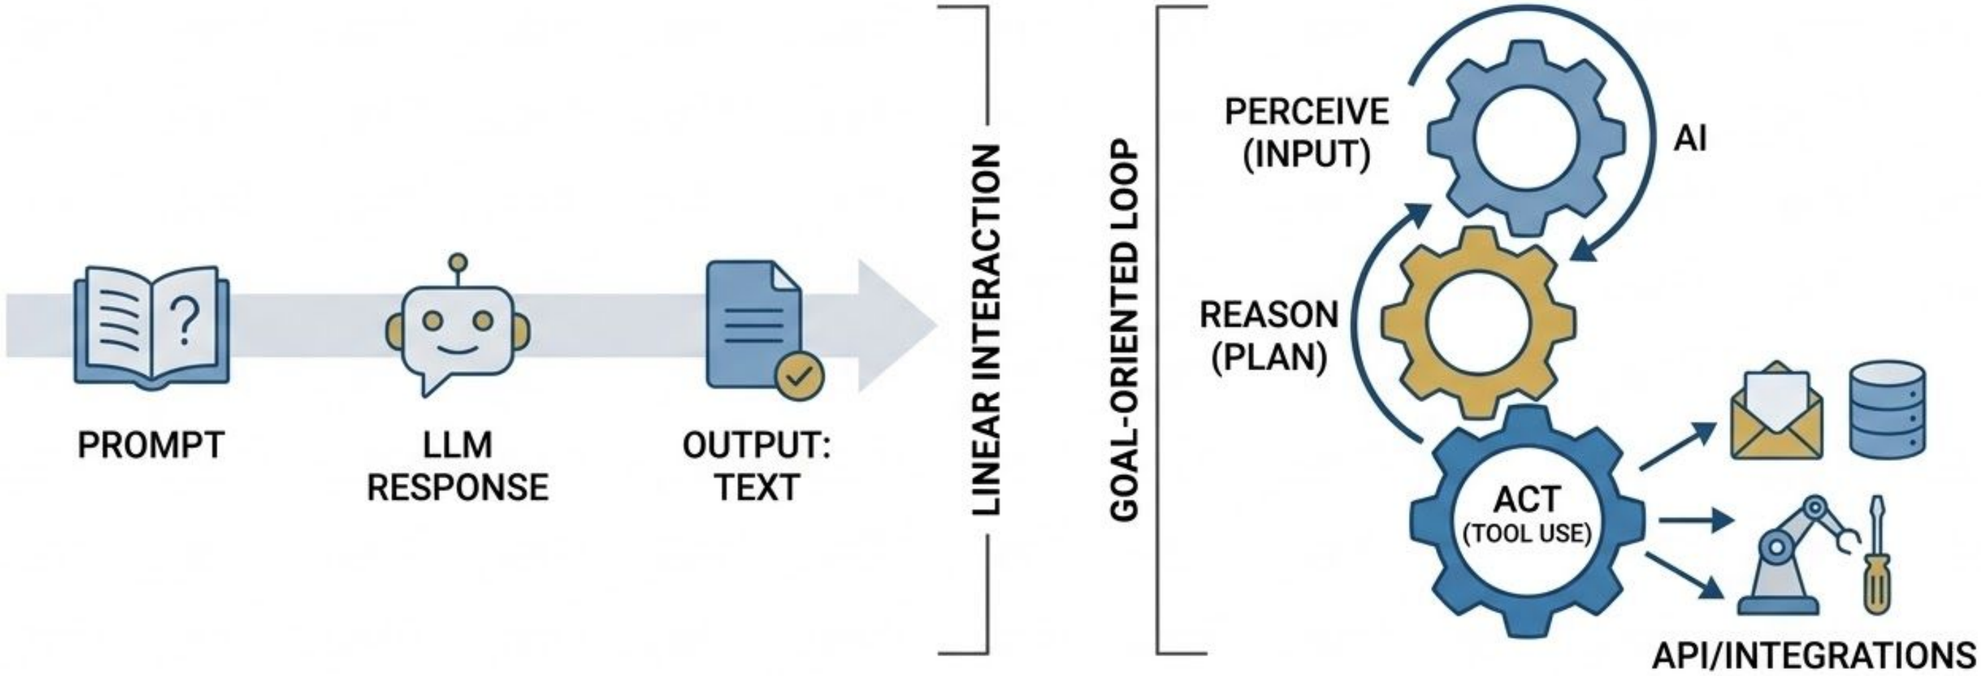

Traditional AI (Linear)

Agentic AI (The Agent)

"By the end of 2026, 40% of enterprise applications will feature task-specific AI agents, up from less than 5% in 2025."

— *Gartner*

Excessive Agency

Goal Hijacking

The Speed of Failure

Are Agentic AI Threats completely new?

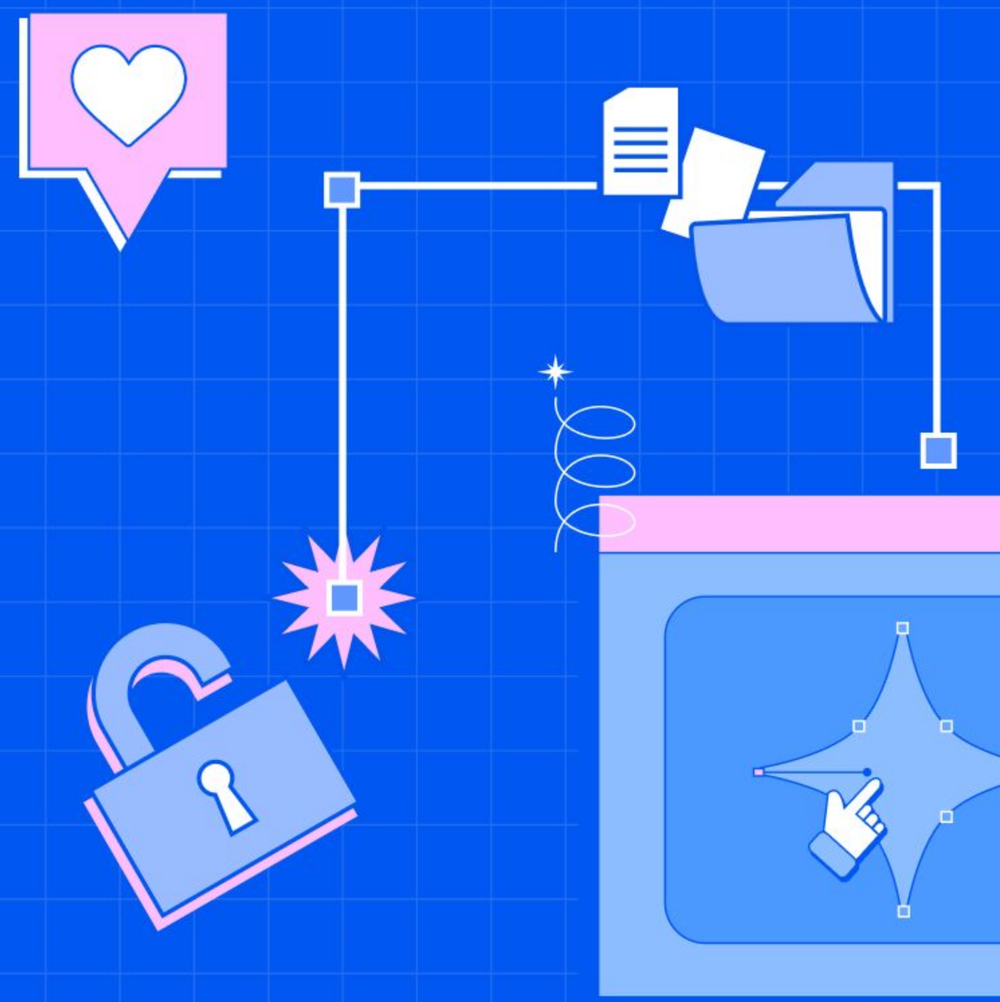

Overview of the AI Stack

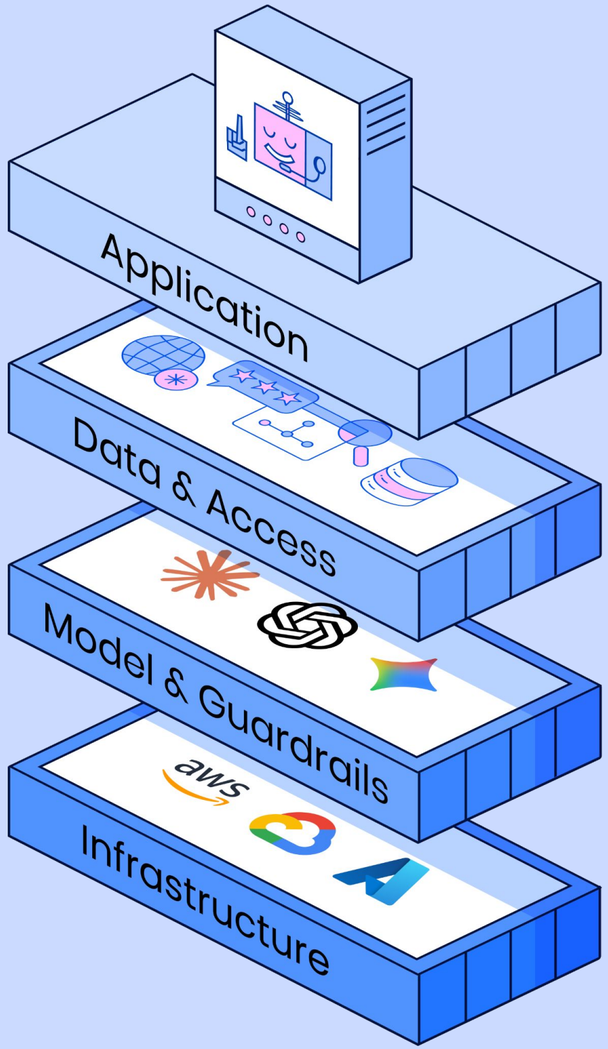

Putting the AI Puzzle Together

1. Infrastructure (CNAPP)
2. Model (AISPM)
3. Data (DSPM)
4. Application (AppSec)

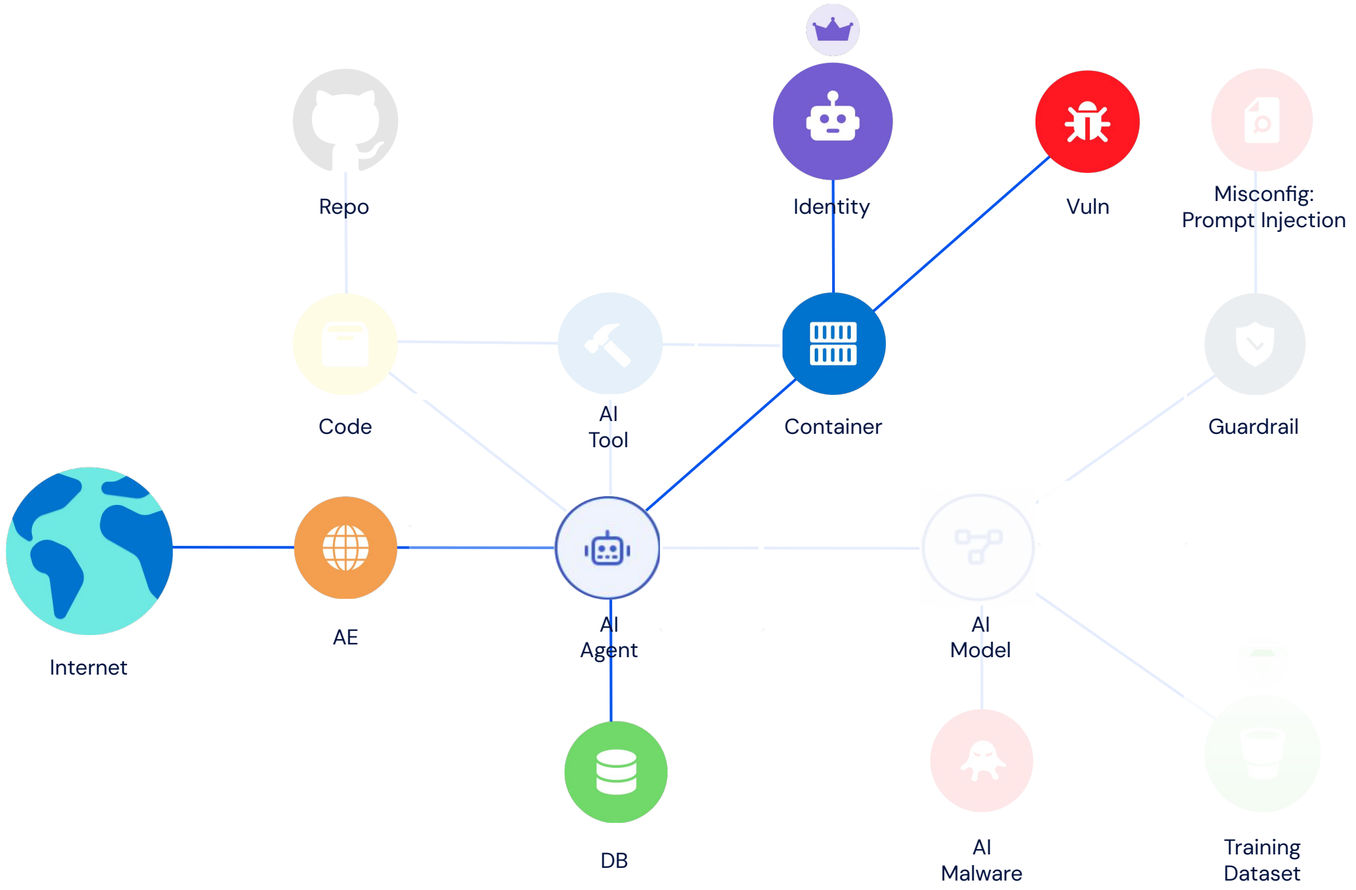

Putting the AI Puzzle Together

1. Infrastructure (CNAPP)
2. Model (AISPM)
3. Data (DSPM)
4. Application (AppSec)

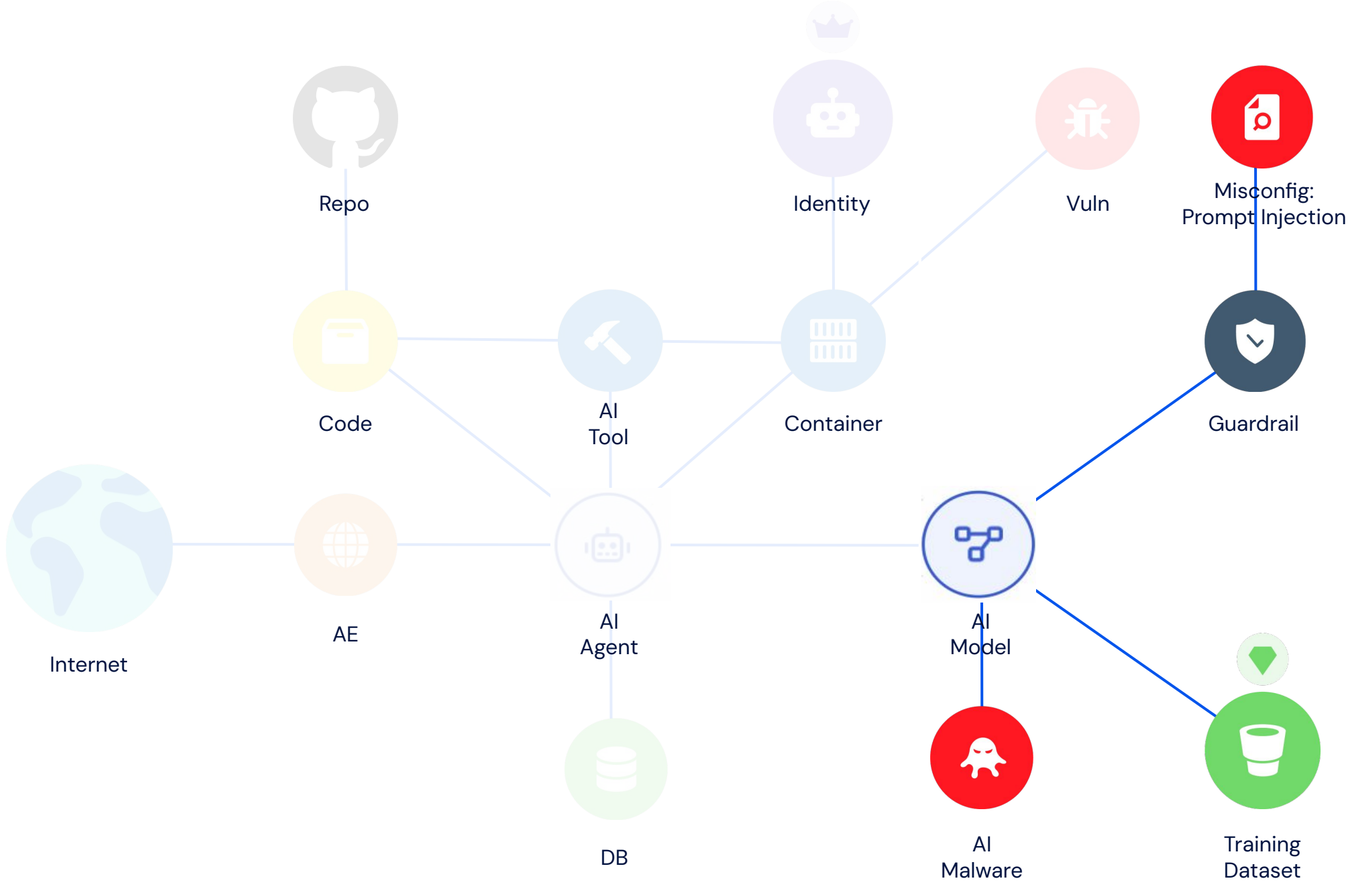

Putting the AI Puzzle Together

1. Infrastructure (CNAPP)
2. Model (AISPM)
3. Data (DSPM)
4. Application (AppSec)

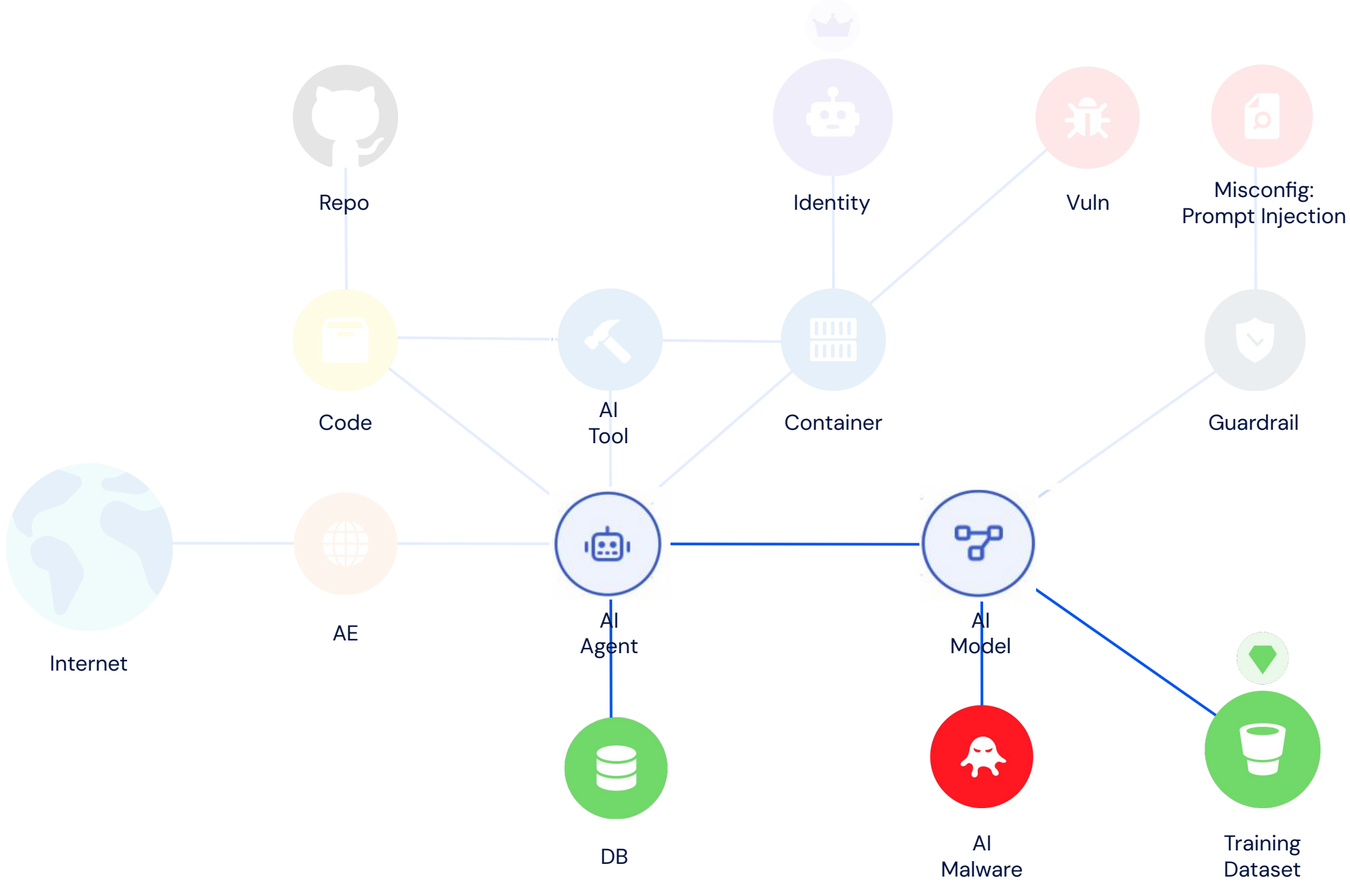

Putting the AI Puzzle Together

1. Infrastructure (CNAPP)
2. Model (AISPM)
3. Data (DSPM)
4. Application (AppSec)

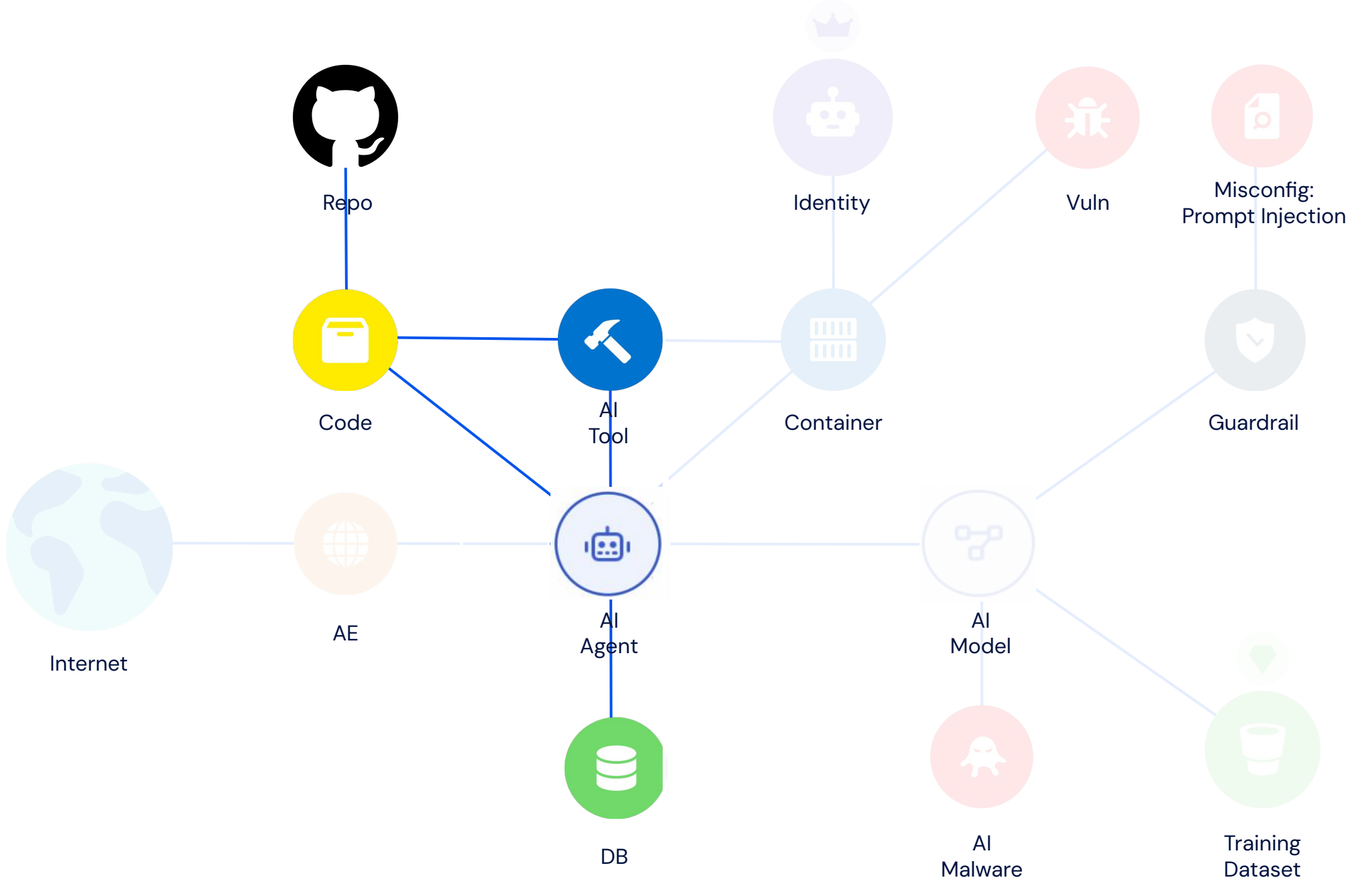

Let’s look at an example

Sensitive data exposure via an externally accessible AI chatbot

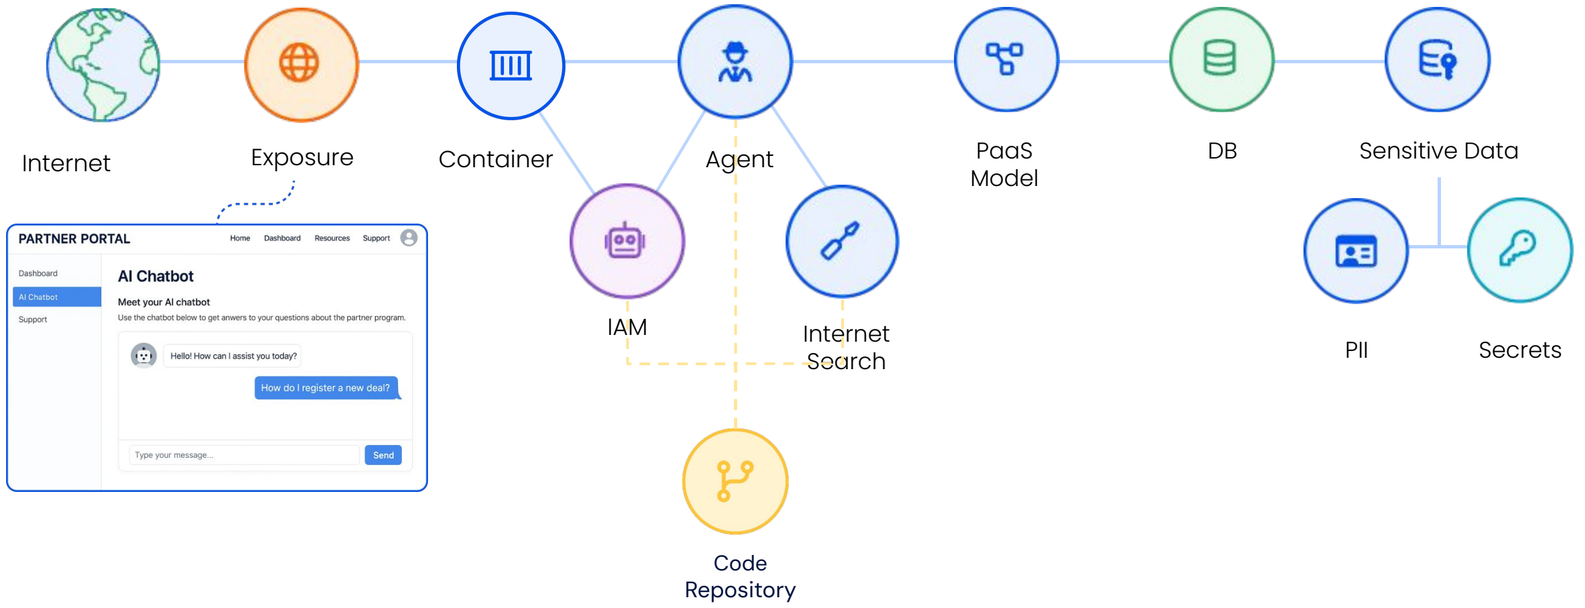

Infra & Access

Hosted AI Agent (chatbot), defined in code, running on container, externally exposed to the internet.

Model & Guardrails

Agent running on PaaS AI Model, was fined tuned on Enterprise data. No Guardrails have been configured.

Data

Training data used for AI model contains PII and Secrets.

Application

Hosted AI Agent has AI Tool capable of searching internet.

Are Agentic AI Threats completely new?

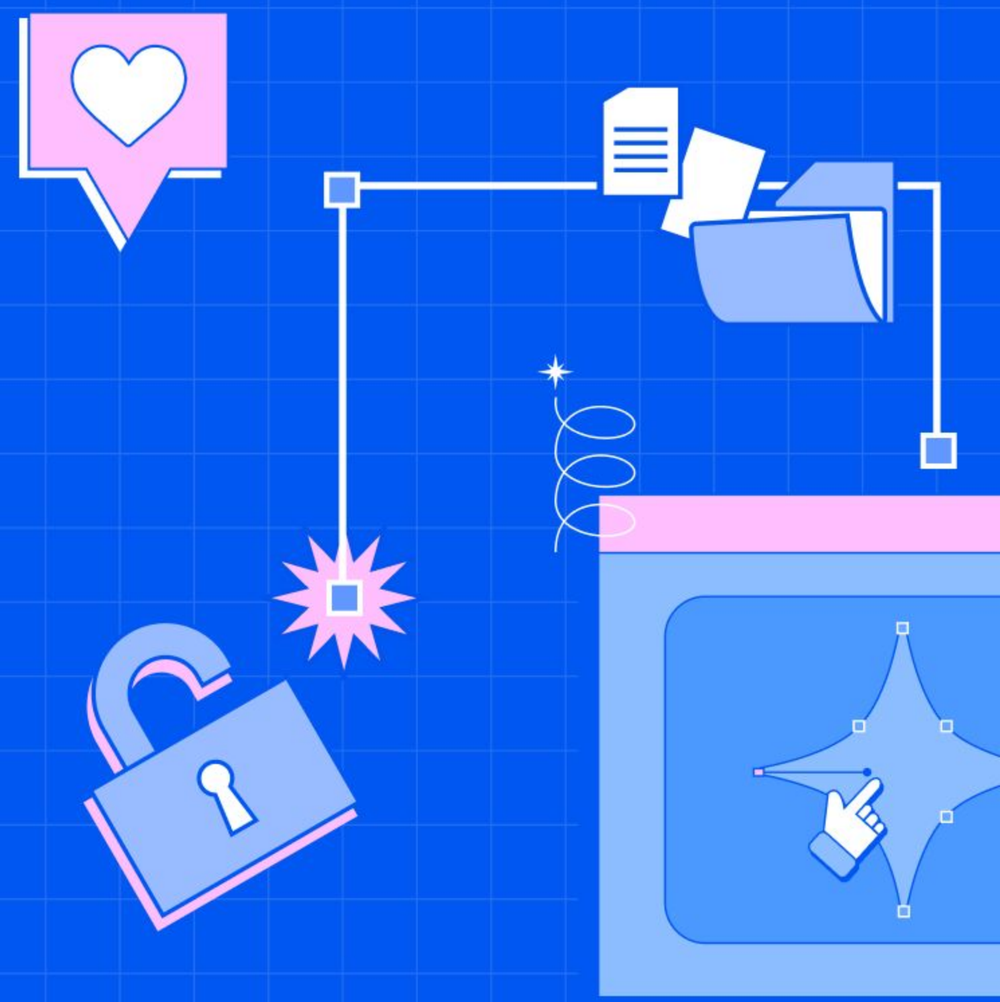

Is the AI Stack All New?

Not New

Infra & Access

AI Workloads & Identities

New

Model & Guardrails

Hosted, PaaS & SaaS

Not New

Data

Training, KB, Customer Data

Kind of New

Application

Agents, Tool, Skills, VCS