CISCO Connect

Digital Resilience in the AI Era

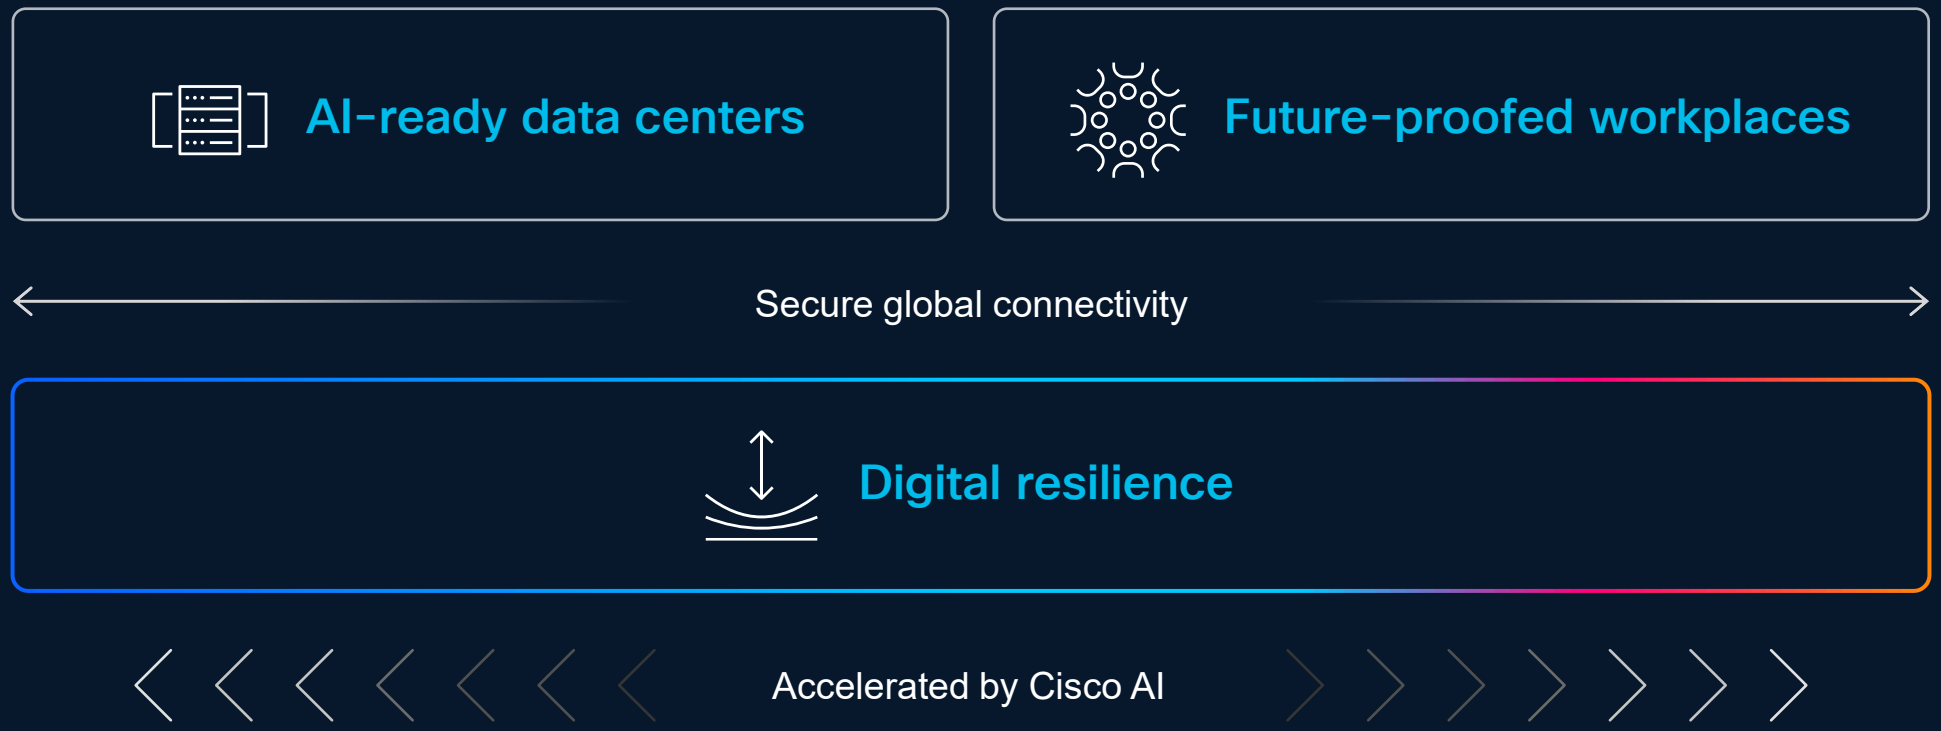
Building a secure and trusted agentic enterprise

Gerry D'Costa
Staff AI Solutions Architect, LinkedIn @gerrydcosta



Date placeholder

Cisco powers how people and technology work together across the physical and digital worlds



Cisco Confidential

Digital Resilience

Being able to adapt to changing risk

What does resilience mean?



The ability to **recover quickly and adapt** as the landscape and risk **changes around you**

Cisco Confidential

Digital Resilience

Being able to adapt to changing risk

from an AI perspective...

AI Digital Resilience

ability to **validate and control**
AI Systems & Agents from going rogue

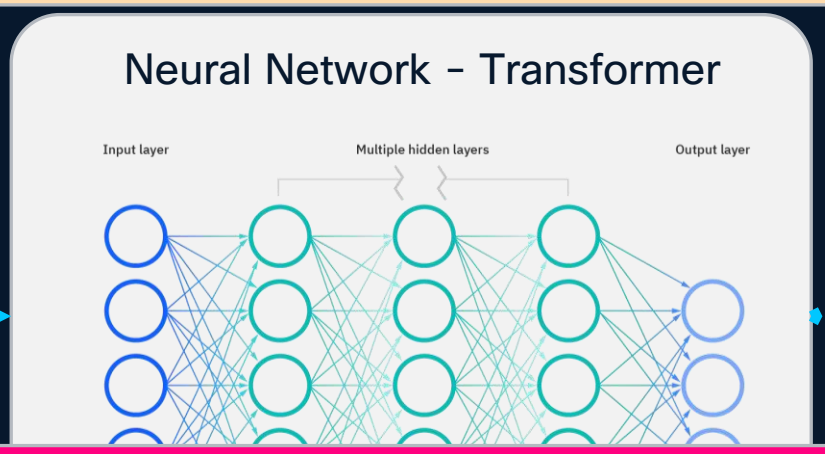
Cisco Confidential

Digital Resilience

Being able to adapt to changing risk

from an AI perspective...

AI Digital Resilience



"tell me a story about a princess"

"once upon a time..."

"long ago and far away..."

non-deterministic reasoning is AI's power to level-up practitioners - **but also, inherently risky**

Cisco Confidential

Digital Resilience

Being able to adapt to changing risk

from an AI perspective...

AI Digital Resilience

ability to **validate and control**
AI Systems & Agents from going rogue

Data Platform built for AI

fast + scalable deterministic data platforms to feed AI systems

AI Governance Visibility

the creation, validation & adherence of governance policies

AI Security & Observability

ability to monitor + enforce runtime security policies & agent behaviour

AI Cost Management

ability to monitor & optimize value-based token consumption

Cisco Confidential

Digital Resilience

Being able to adapt to changing risk

Data Platform built for AI

fast + scalable deterministic data platforms to feed AI systems

AI Governance Visibility

the creation, validation & adherence of governance policies

AI Security & Observability

ability to monitor + enforce runtime security policies & agent behaviour

AI Cost Management

ability to monitor & optimize value-based token consumption

The Foundation of AI Digital Resilience is achieved through

AI Trust & Assurance

Cisco Confidential

AI Trust & Assurance

needs a fast & scalable data platform

Data Platform built for AI

fast + scalable deterministic data platforms to feed AI systems

AI Governance Visibility

the creation, validation & adherence of governance policies

AI Security & Observability

ability to monitor & enforce runtime security policies & agent behaviour

AI Cost Management

ability to monitor & optimize value-based token consumption

Cisco Data Fabric – Powered by Splunk

The foundation for Agentic AI

Cisco Data Fabric



Splunk Platform

NEW



Machine Data Lake
parquet format

ludicrous scale

Cost efficient

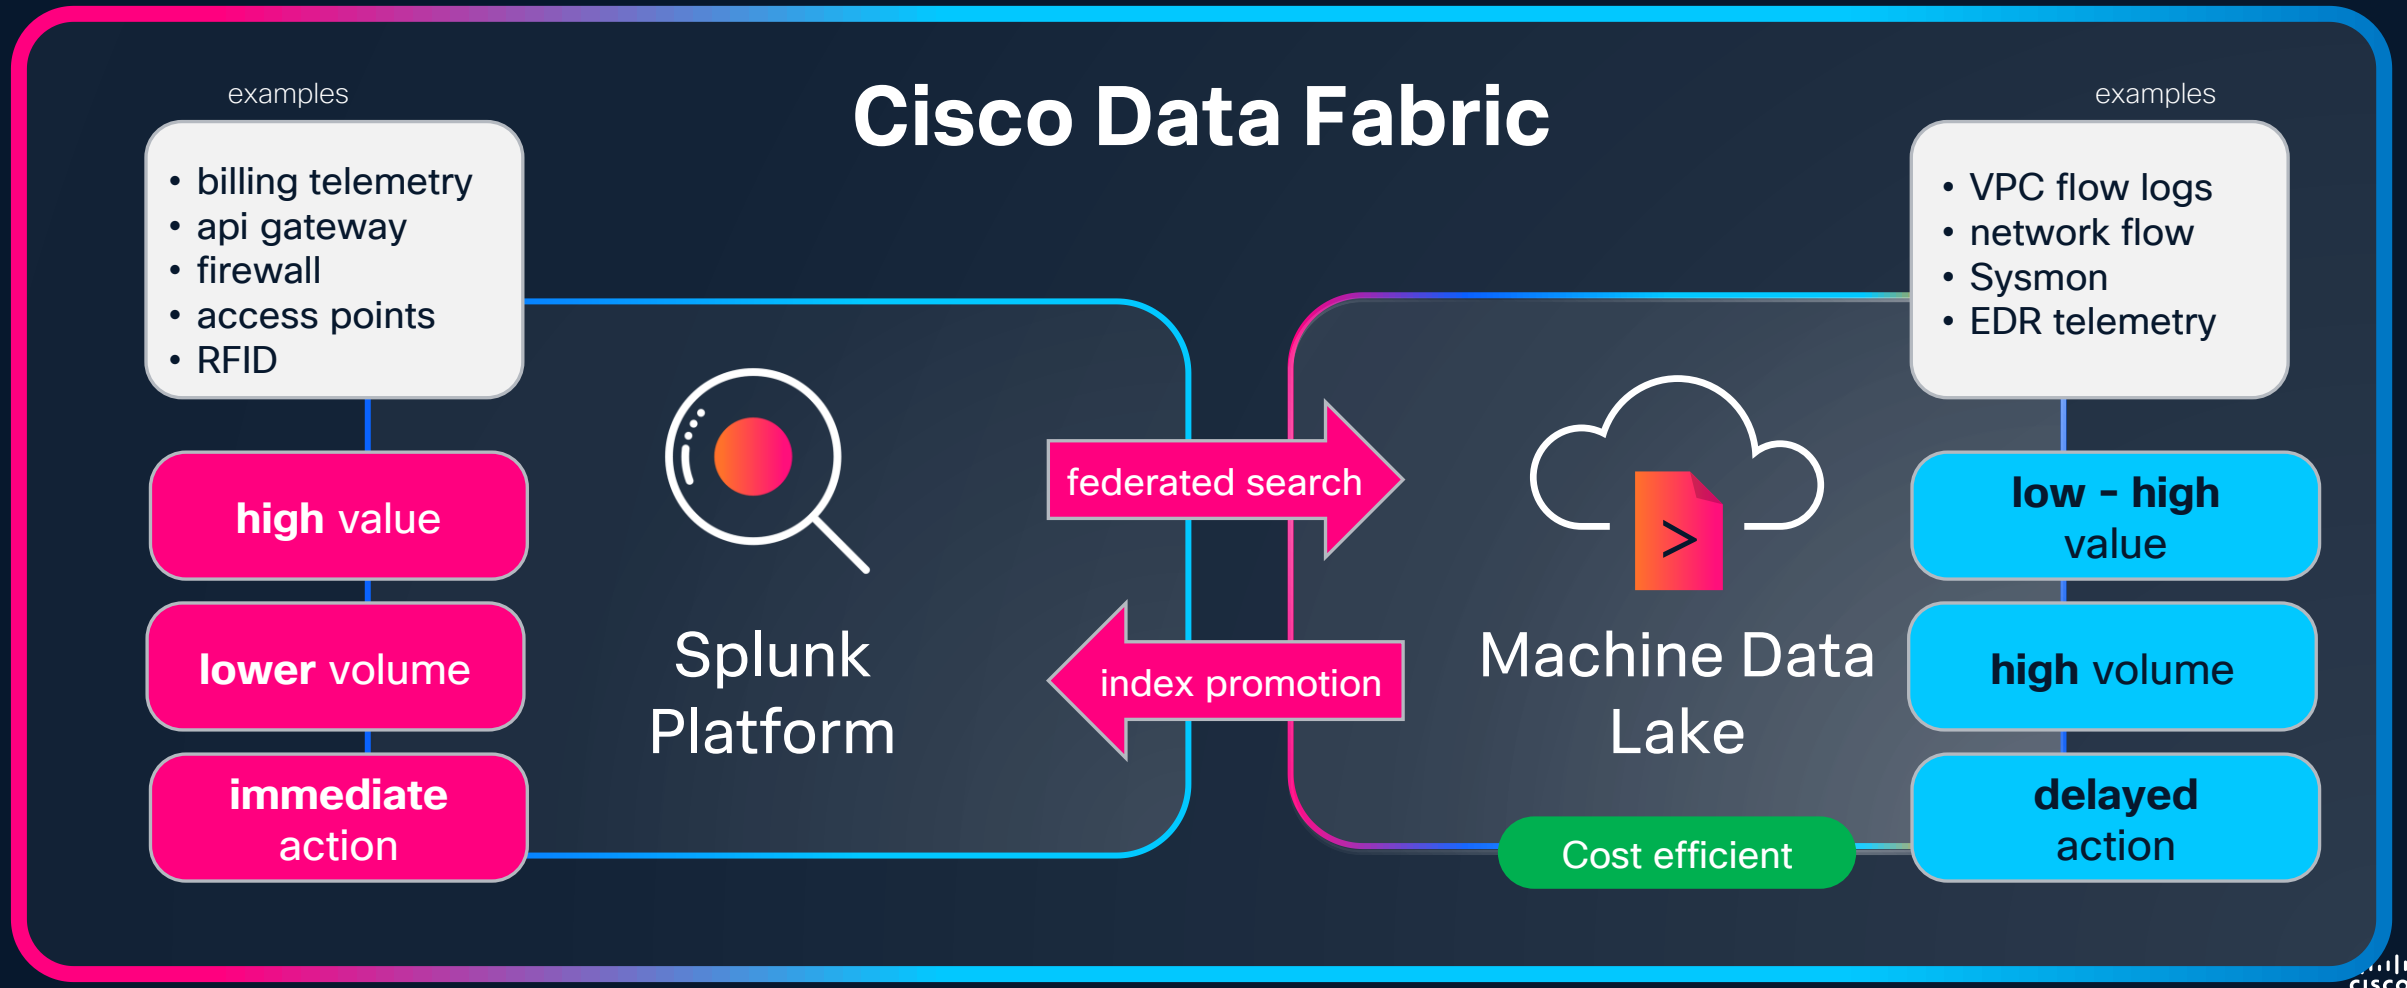


AI-powered Splunk data management

Cisco Confidential

Cisco Data Fabric – Powered by Splunk

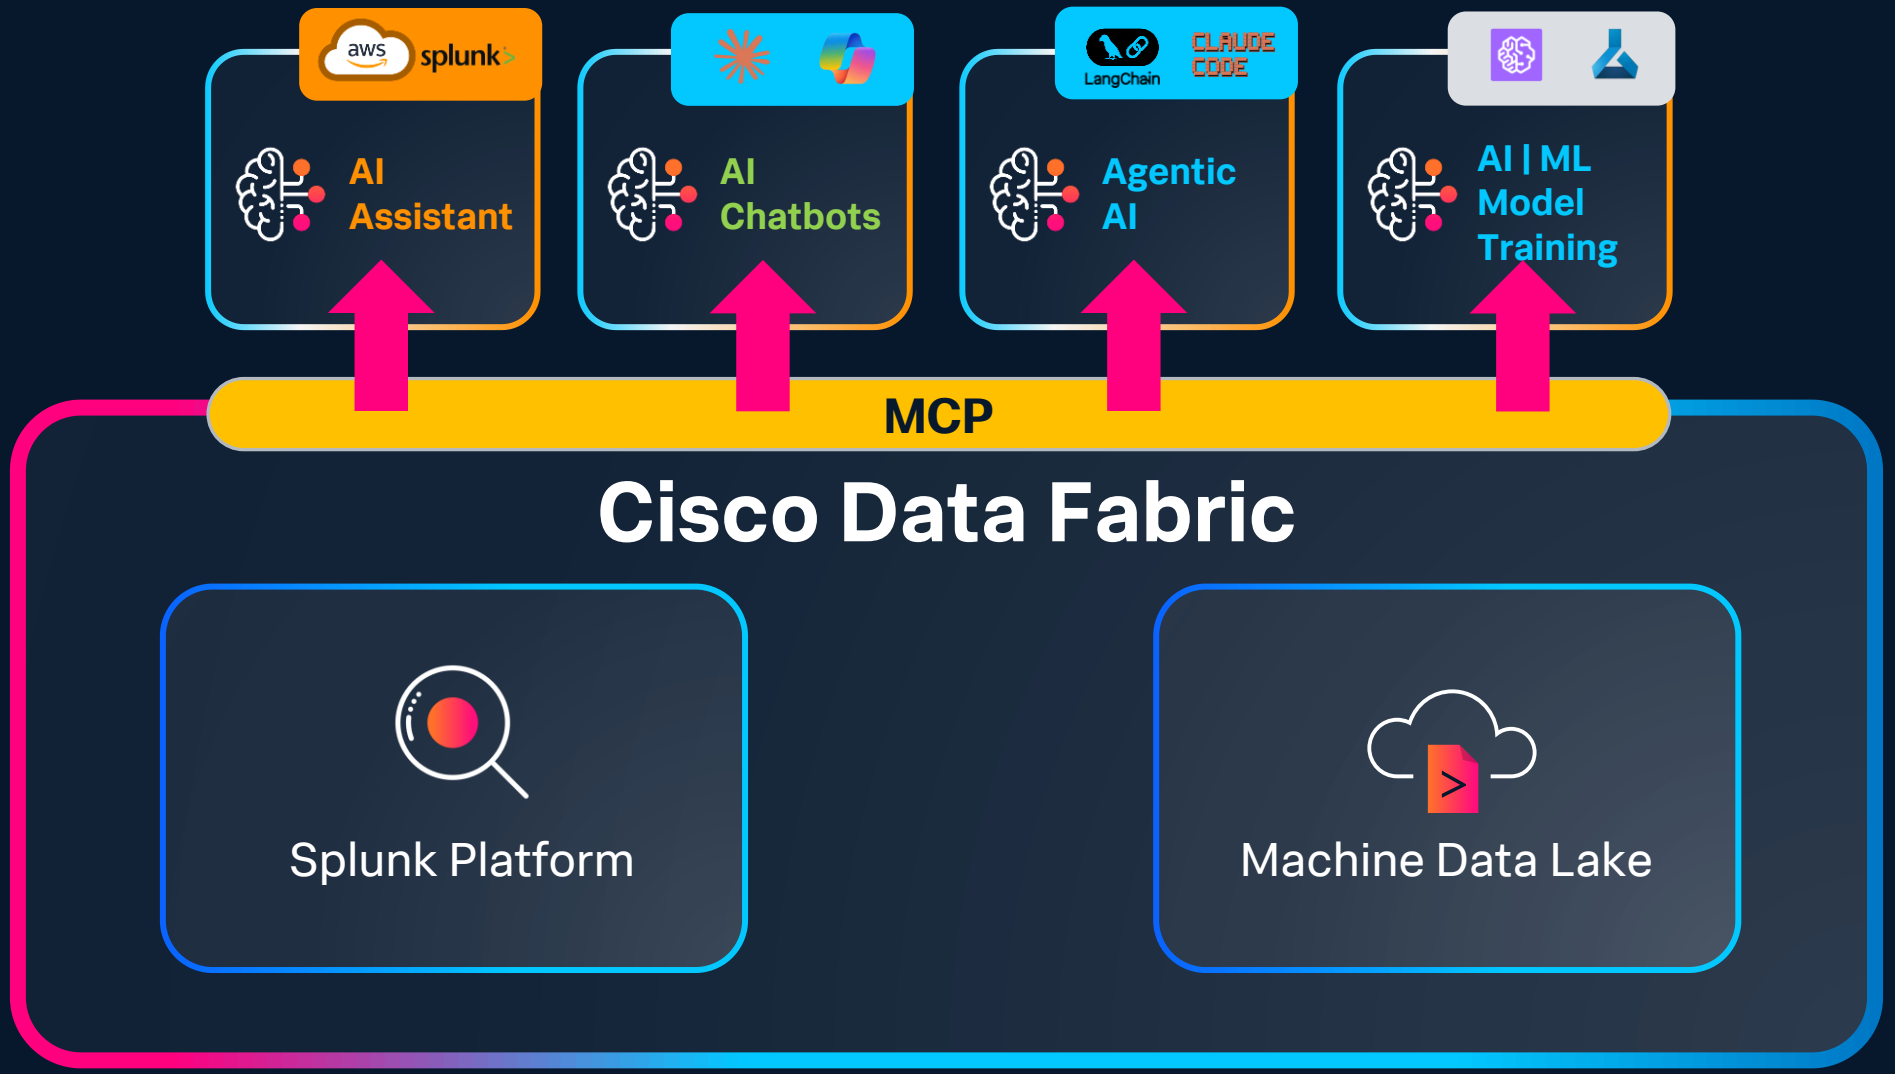
Splunk Cloud Platform vs. Machine Data Lake



Cisco Confidential

Cisco Data Fabric – Powered by Splunk

Enables Agentic AI



Cisco Confidential

AI Trust & Assurance

needs AI governance policies

Data Platform built for AI

fast + scalable deterministic data platforms to feed AI systems

AI Governance Visibility

the creation, validation & adherence of governance policies

AI Security & Observability

ability to monitor & enforce runtime security policies & agent behaviour

AI Cost Management

ability to monitor & optimize value-based token consumption

Building an AI Program

Required business groups



AI Steering Committee

Identify AI Governance Requirements

- What are we authorizing & what does success look like?
- Why do we need AI, rather than conventional automation?
- Defining risk, budget & cost thresholds
- What is required adherence to data compliance policies?
- Public communication position if an AI Agent goes rogue?



AI Centre of Excellence

AI Frameworks, Monitoring, Tools & Enforcement

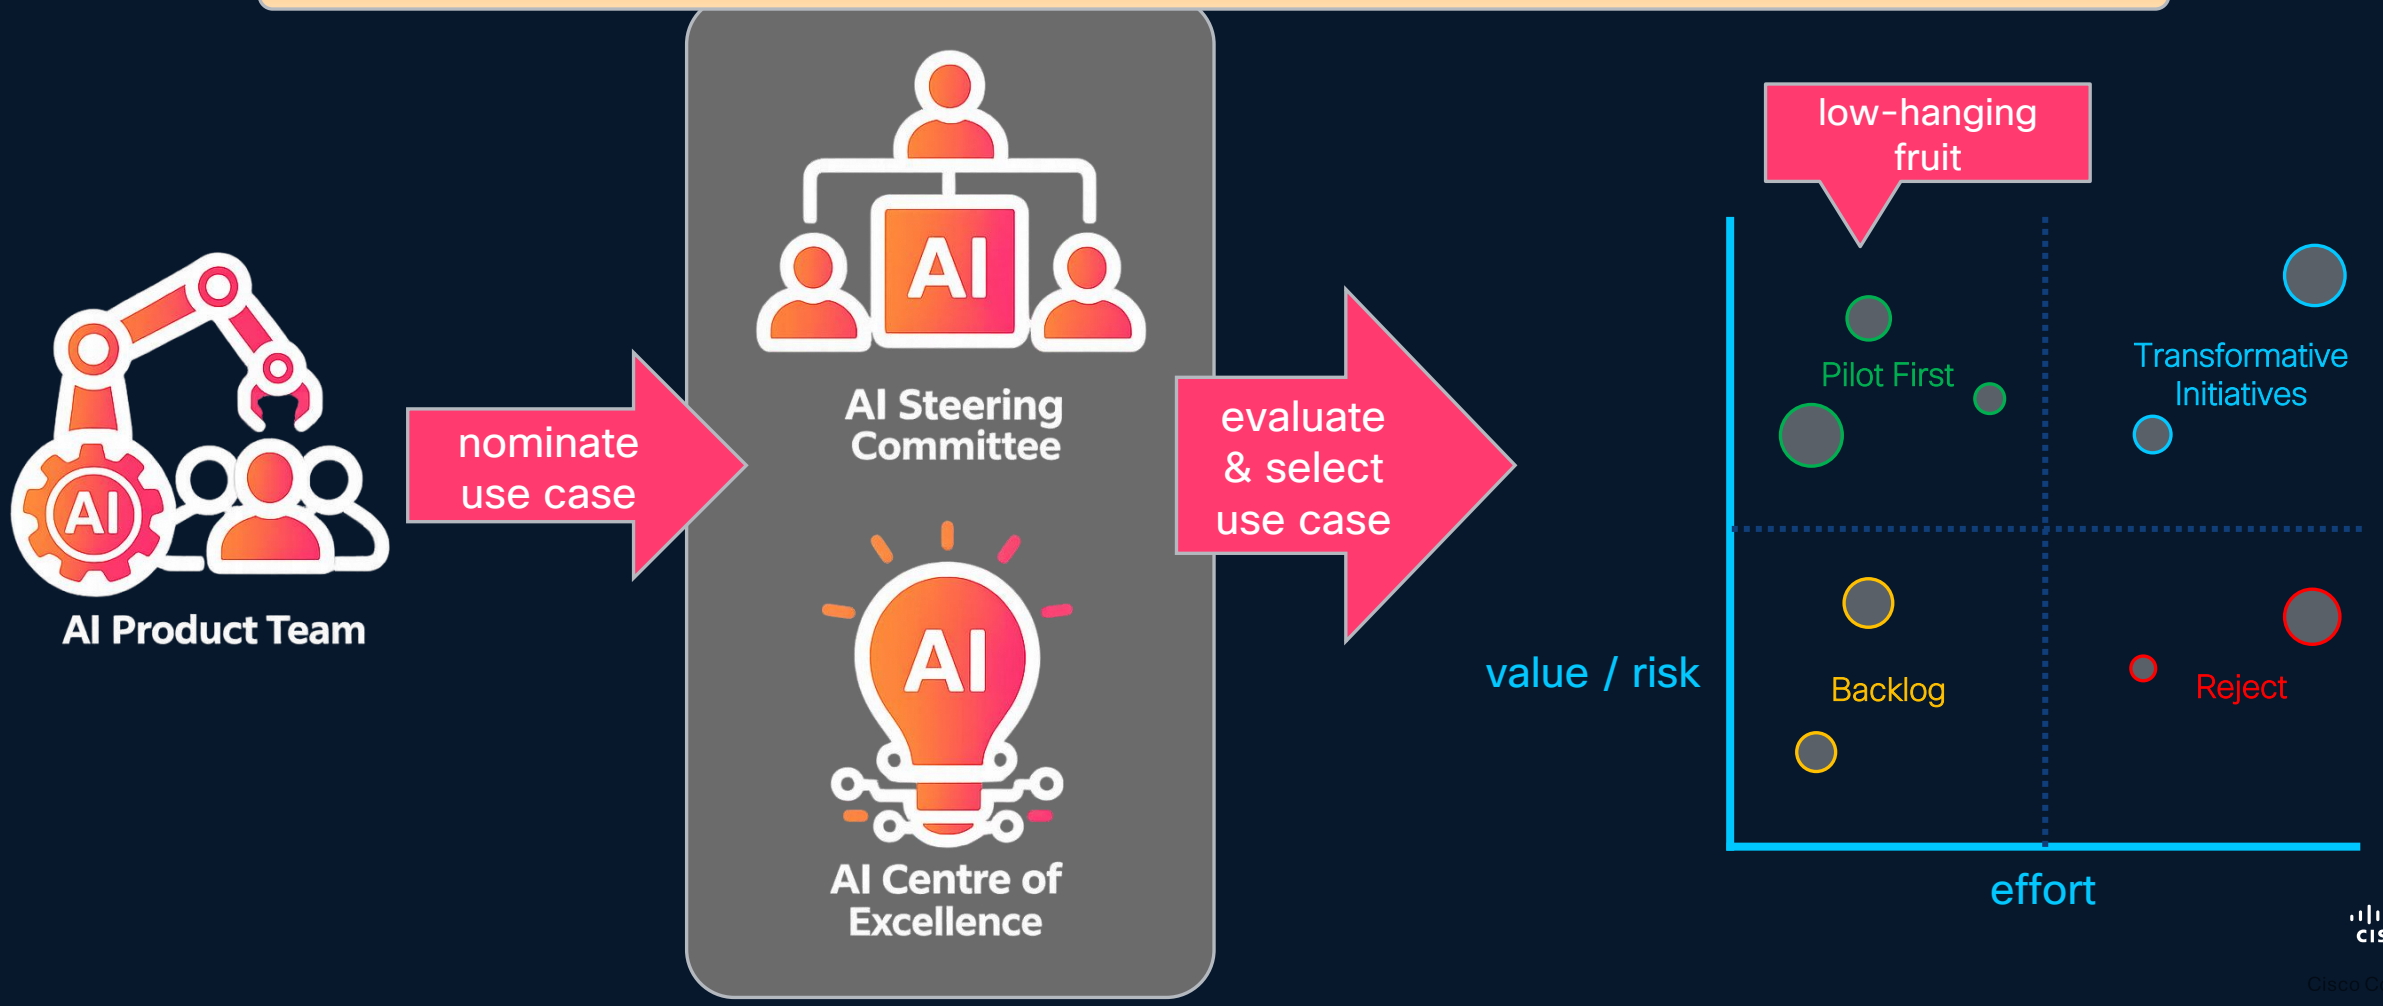
- What is our AI Framework?
- What do we do if we exceed cost thresholds? *Throttle? Alert? Hard-stop?*
- What tools do we use for layered security?
- What tools do we use to monitor & verify AI security adherence and agent behaviour?

Cisco Confidential

Building a minimum viable AI Program (MVAP)

Prioritizing Use Cases

build a minimum viable program that **delivers one secure, cost-controlled outcome,** then expand



Cisco Confidential

Delivering AI Use Cases

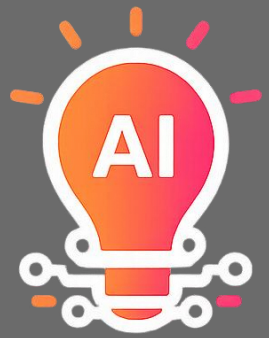
Once identified, each group participates in use case delivery

deliver on use case



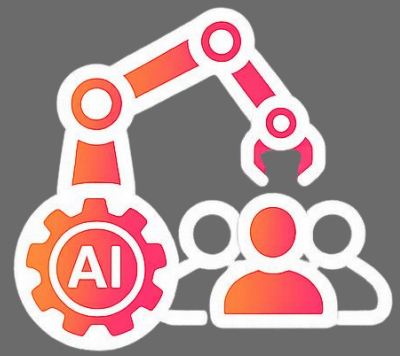
AI Steering Committee

Identify AI Governance Requirements



AI Centre of Excellence

Frameworks, Tools & Enforcement methods



AI Product Team

Build, Enforce and Verify Agentic Systems using standards

Cisco Confidential

AI Trust & Assurance

needs layers of enforcement, cost controls and AI observability

Data Platform built for AI

fast + scalable deterministic data platforms to feed AI systems

AI Governance Visibility

the creation, validation & adherence of governance policies

AI Security & Observability

ability to monitor & enforce runtime security policies & agent behaviour

AI Cost Management

ability to monitor & optimize value-based token consumption

Tiers of Agentic Security & Governance Controls

When developing deterministic apps, control was generally clear

General rules for traditional app control

The app **can** do this

The app **can't do** that

For AI, the controls are more complicated

Cisco Confidential

Tiers of Agentic Security & Governance Controls

We need **rules** to enforce when we talk about agentic control

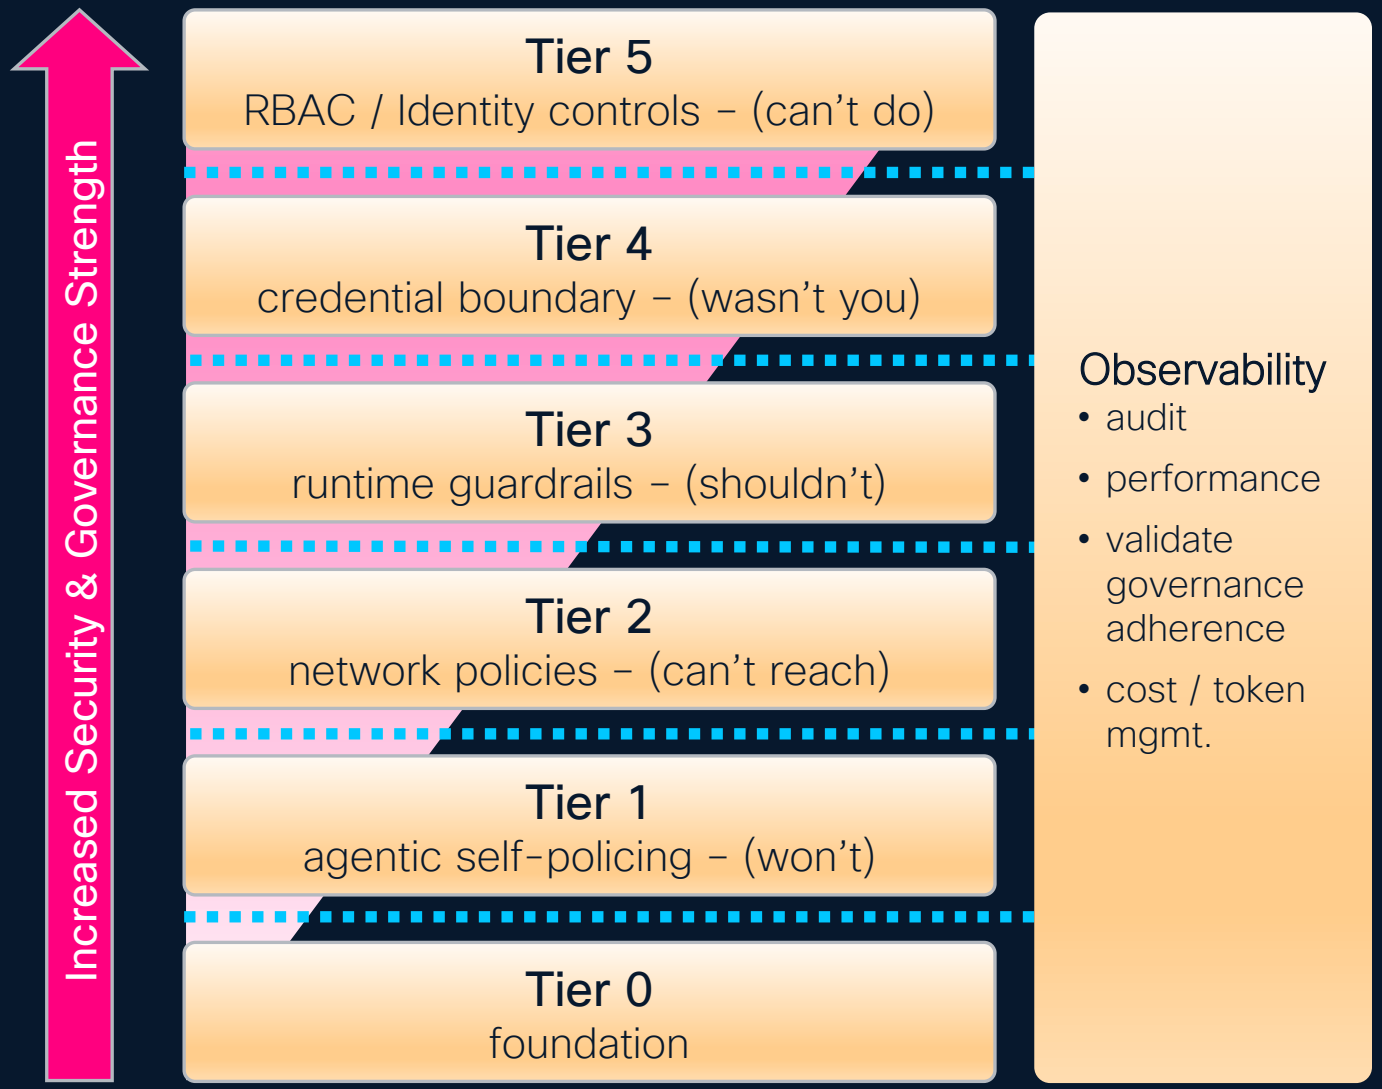
Rules for Agentic Control

- The agentic system **can't do** these things
- That **wasn't** the agentic system accessing data or tools – it should be a proxy system
- The agentic system **shouldn't** be acting on this request
- The agentic system **can't reach** these things
- The agentic system **won't** do these things

Cisco Confidential

Tiers of Agentic Security & Governance Controls

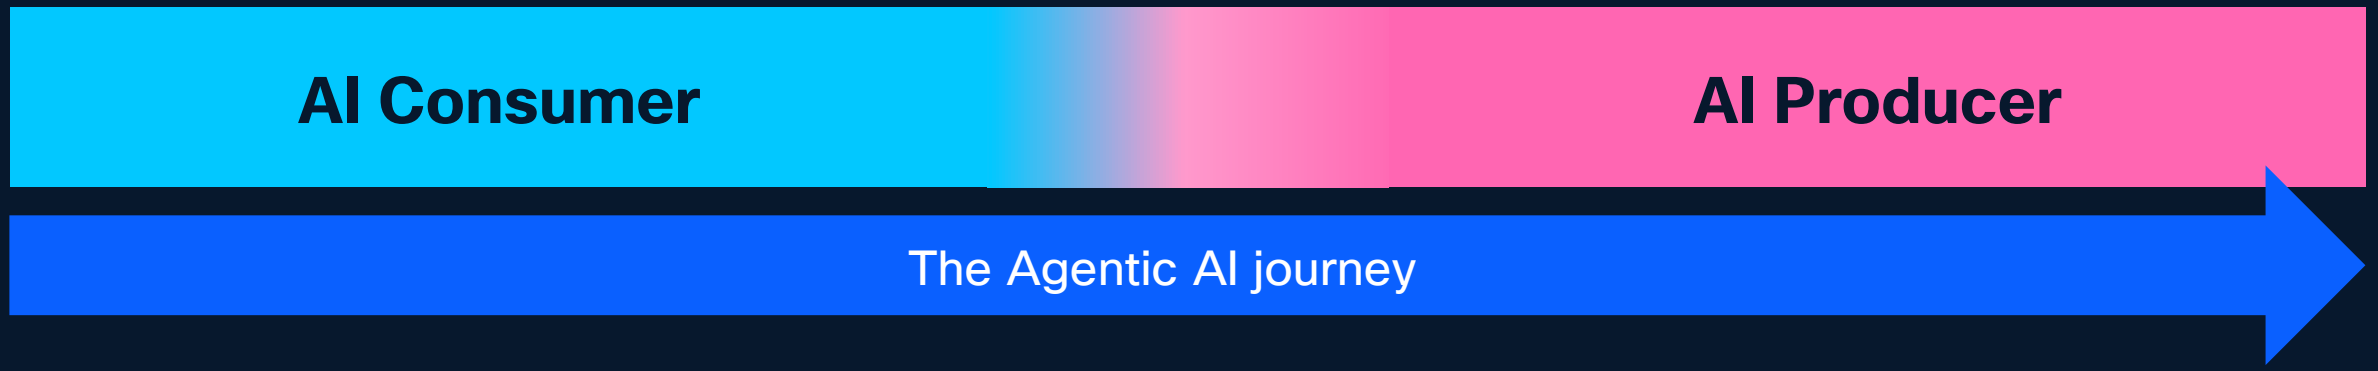
Protection of Agentic systems requires layers of security & controls



Cisco Confidential

Tiers of Agentic Security & Governance Controls

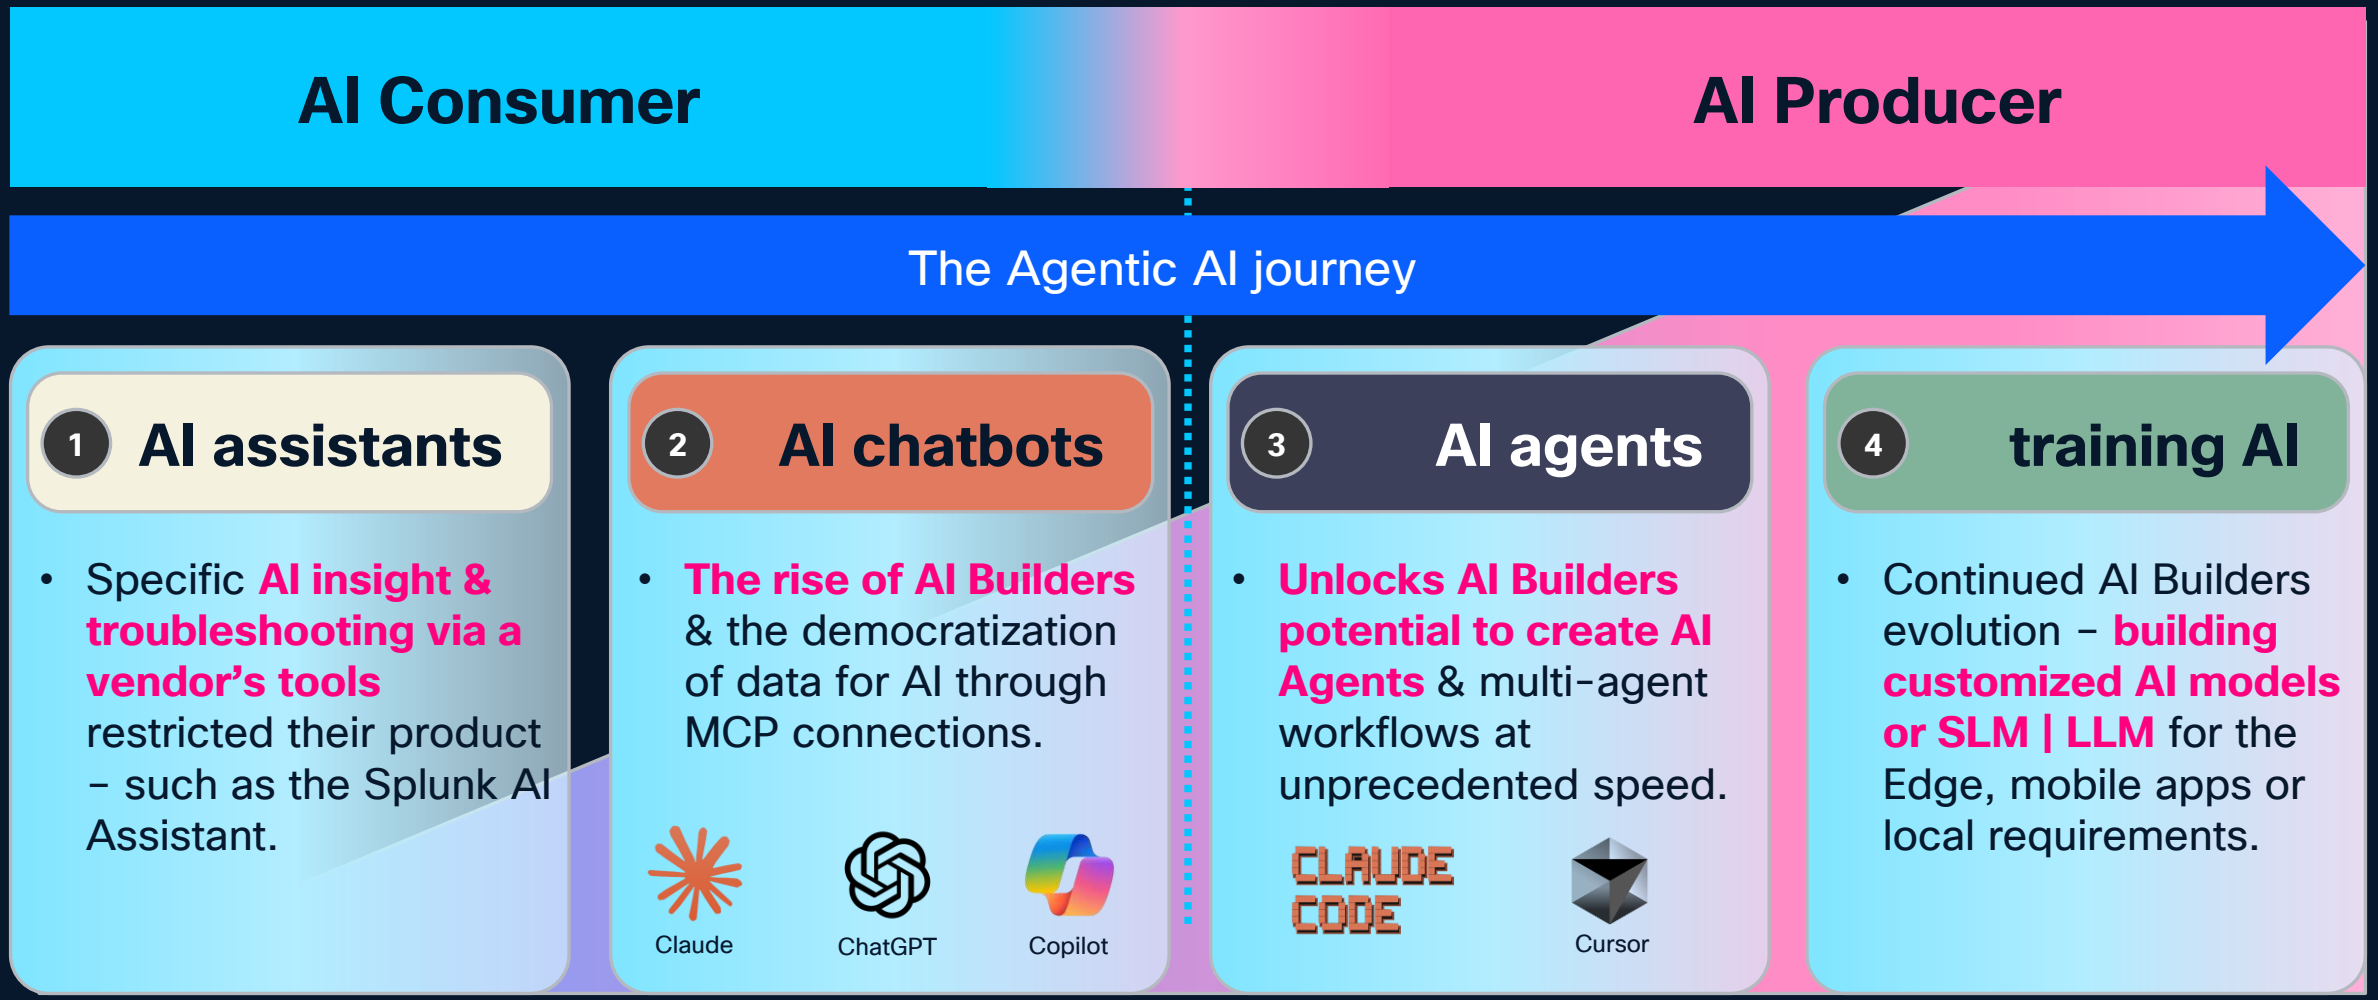
Differs depending on where you are on your journey



Cisco Confidential

Tiers of Agentic Security & Governance Controls

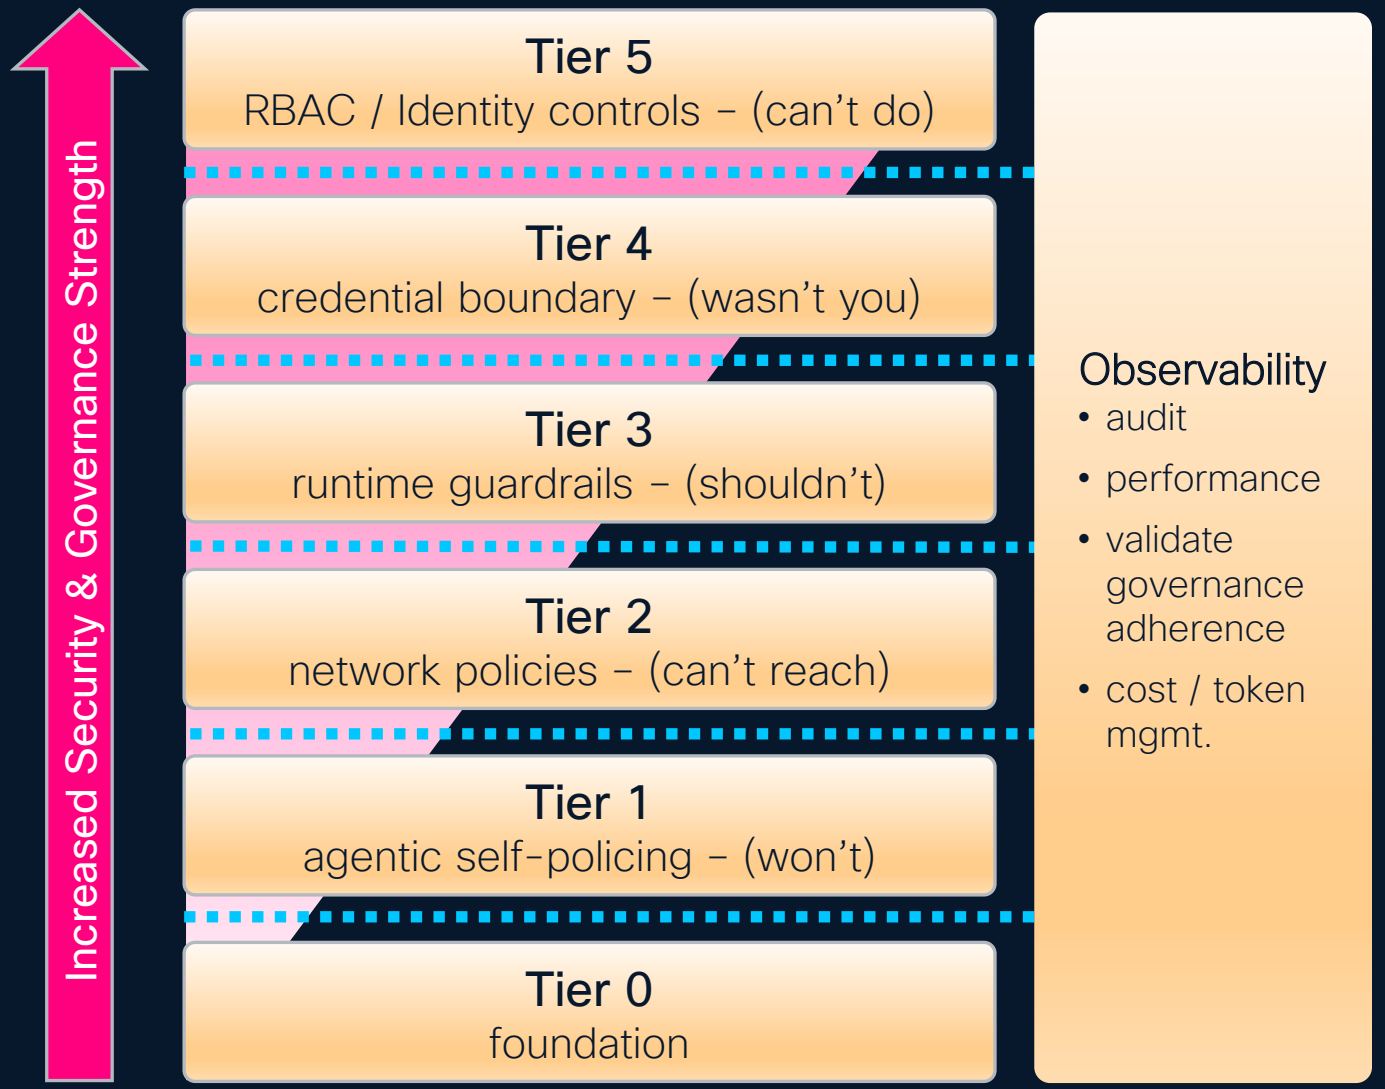
Differs depending on where you are on your journey



Cisco Confidential

Tiers of Agentic Security & Governance Controls

layers of security & controls different depending on adoption journey



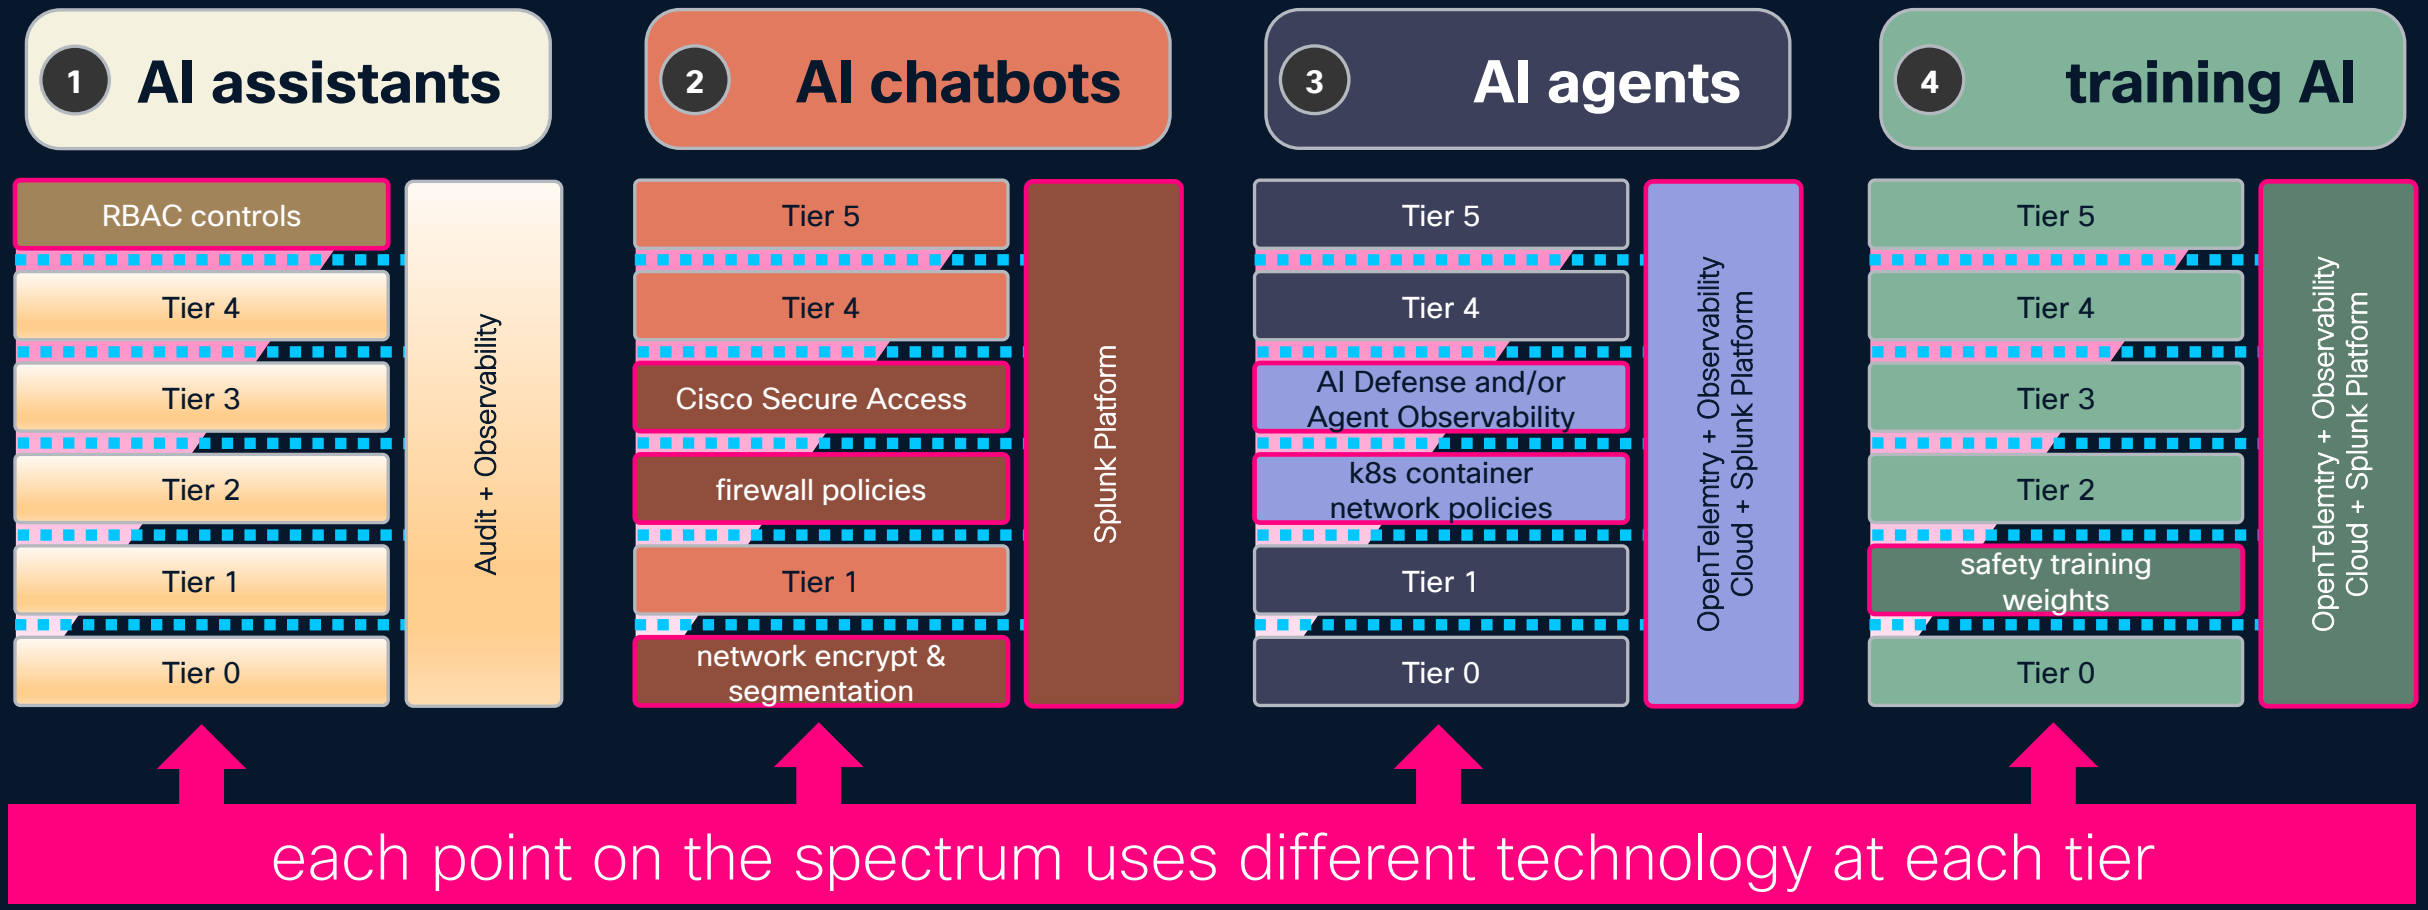
align differently

1. AI assistants
2. AI chatbots
3. AI agents
4. training AI

Cisco Confidential

Tiers of Agentic Security & Governance Controls

Protection of Agentic systems requires layers of security & controls

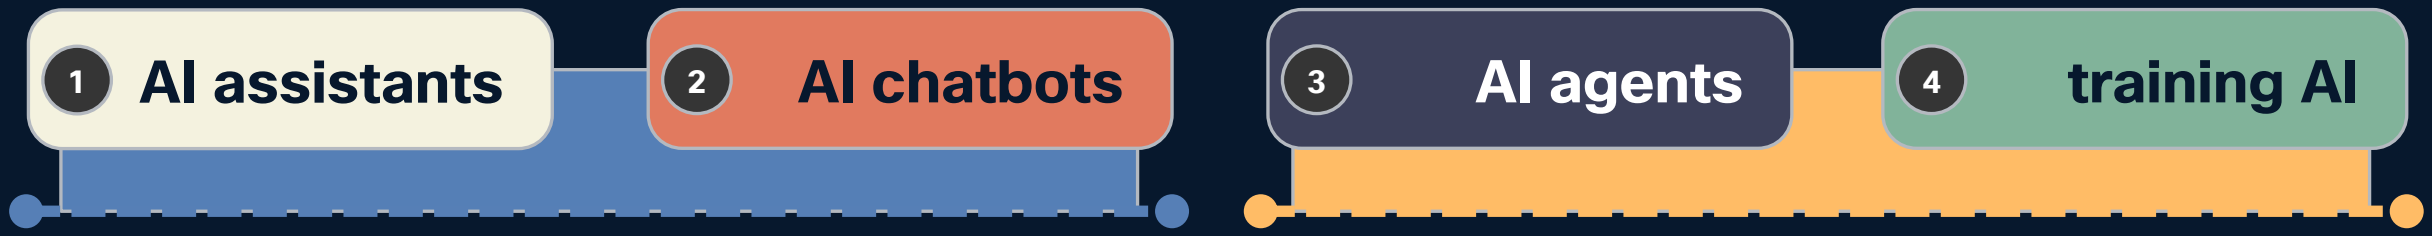


each point on the spectrum uses different technology at each tier

Cisco Confidential

Agentic Monitoring & Observability

1st party AI vs. 3rd party AI – each have different monitoring & observability solutions



3rd party
AI consumption

AI Reasoning occurs **outside** the control of an organization

Ie. Using a frontier model chatbot or AI embedded in a product

1st party
AI consumption

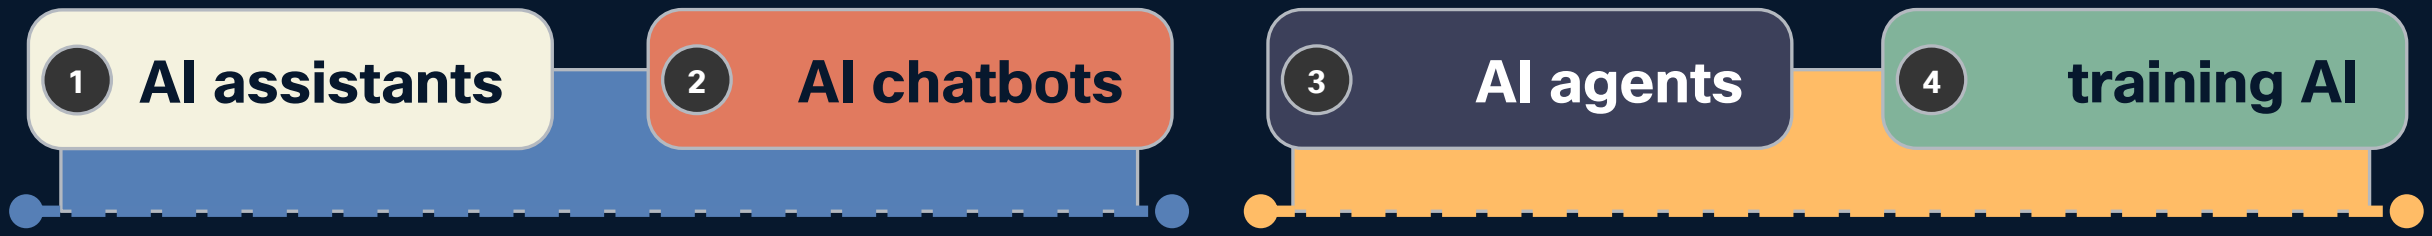
AI Reasoning occurs **within** the control of an organization

Ie. Having build an agentic system where you can control execution

Cisco Confidential

Agentic Monitoring & Observability

1st party AI vs. 3rd party AI – each have different monitoring & observability solutions



3rd party
AI consumption

- Agentic SOC *(using Splunk Enterprise Security)*
- Cisco Secure Access
- Cisco Cloud Control + AI Canvas
- Splunk MCP Server
- Splunk AI Assistant
- Splunk Agent LaunchPad

1st party
AI consumption

- Cisco AI Defense
- Splunk Agent Observability *(formerly Galileo)*
- Cisco AI DefenseClaw
- Splunk Observability Cloud

Cisco Data Fabric

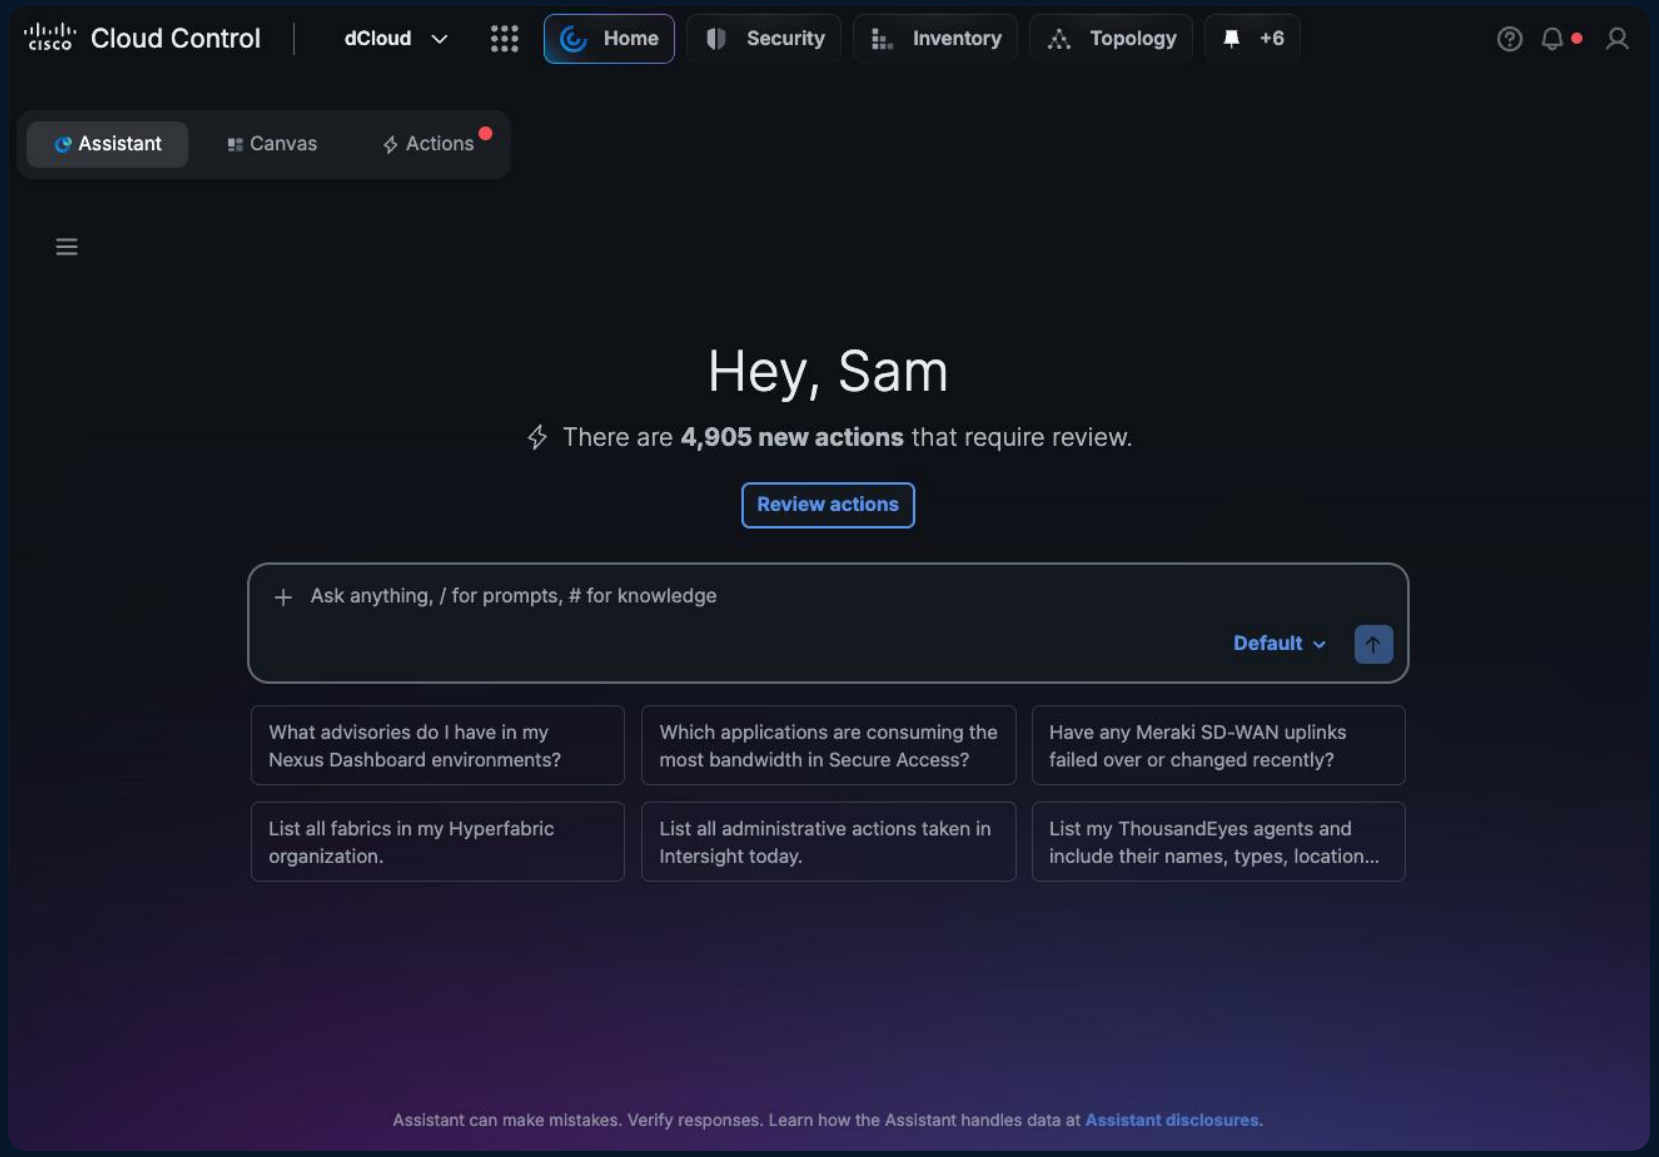
Cisco Confidential

Cisco Cloud Control

Gateway to everything

how can I help you?

- Unified cloud operations platform management environment using Agentic AI combining all Cisco technology – including Splunk.
- Powered by the Cisco Data Fabric.



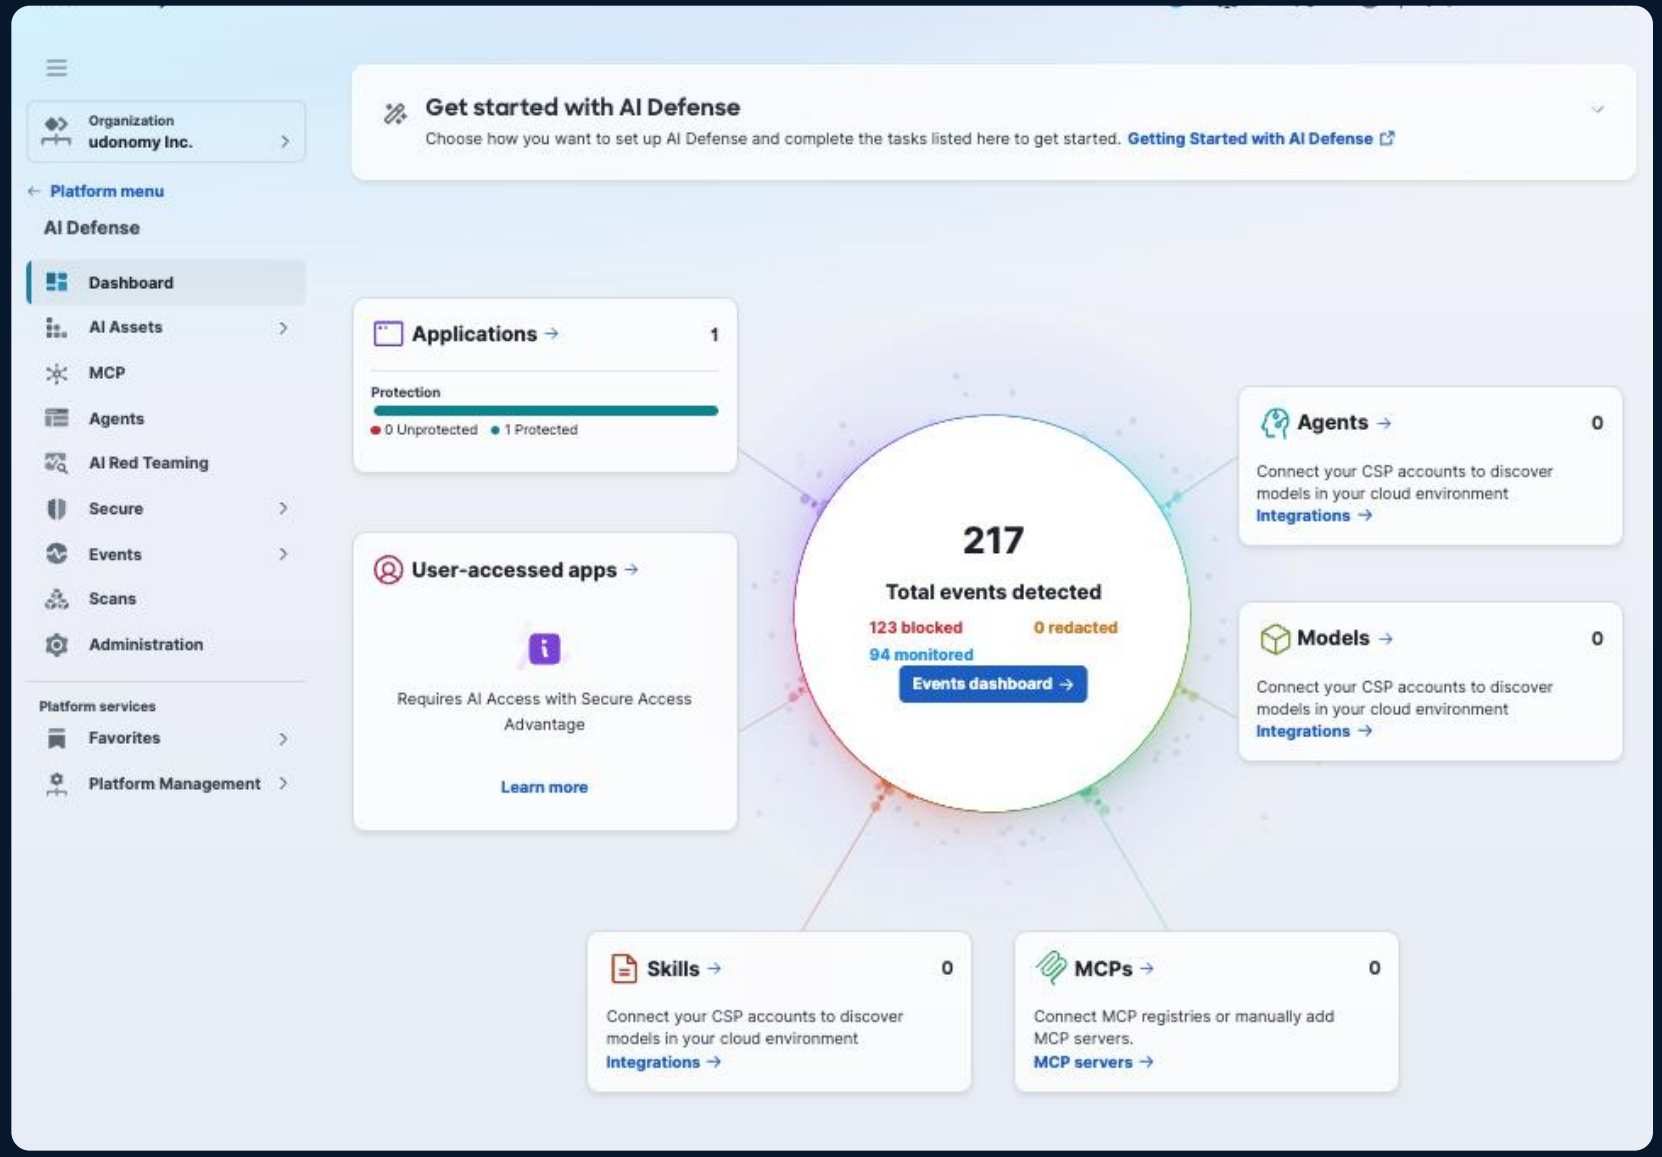
Cisco Confidential

Cisco AI Defense / DefenseClaw

Security Guardrails

is it safe?

- Protects AI Applications / Models / Agents that organizations build.
- Enforcement of runtime guardrails to protect against prompt injection and other agent attacks
- DefenseClaw can be wired into the scaffold of the agent for low-latency or air-gapped protection.

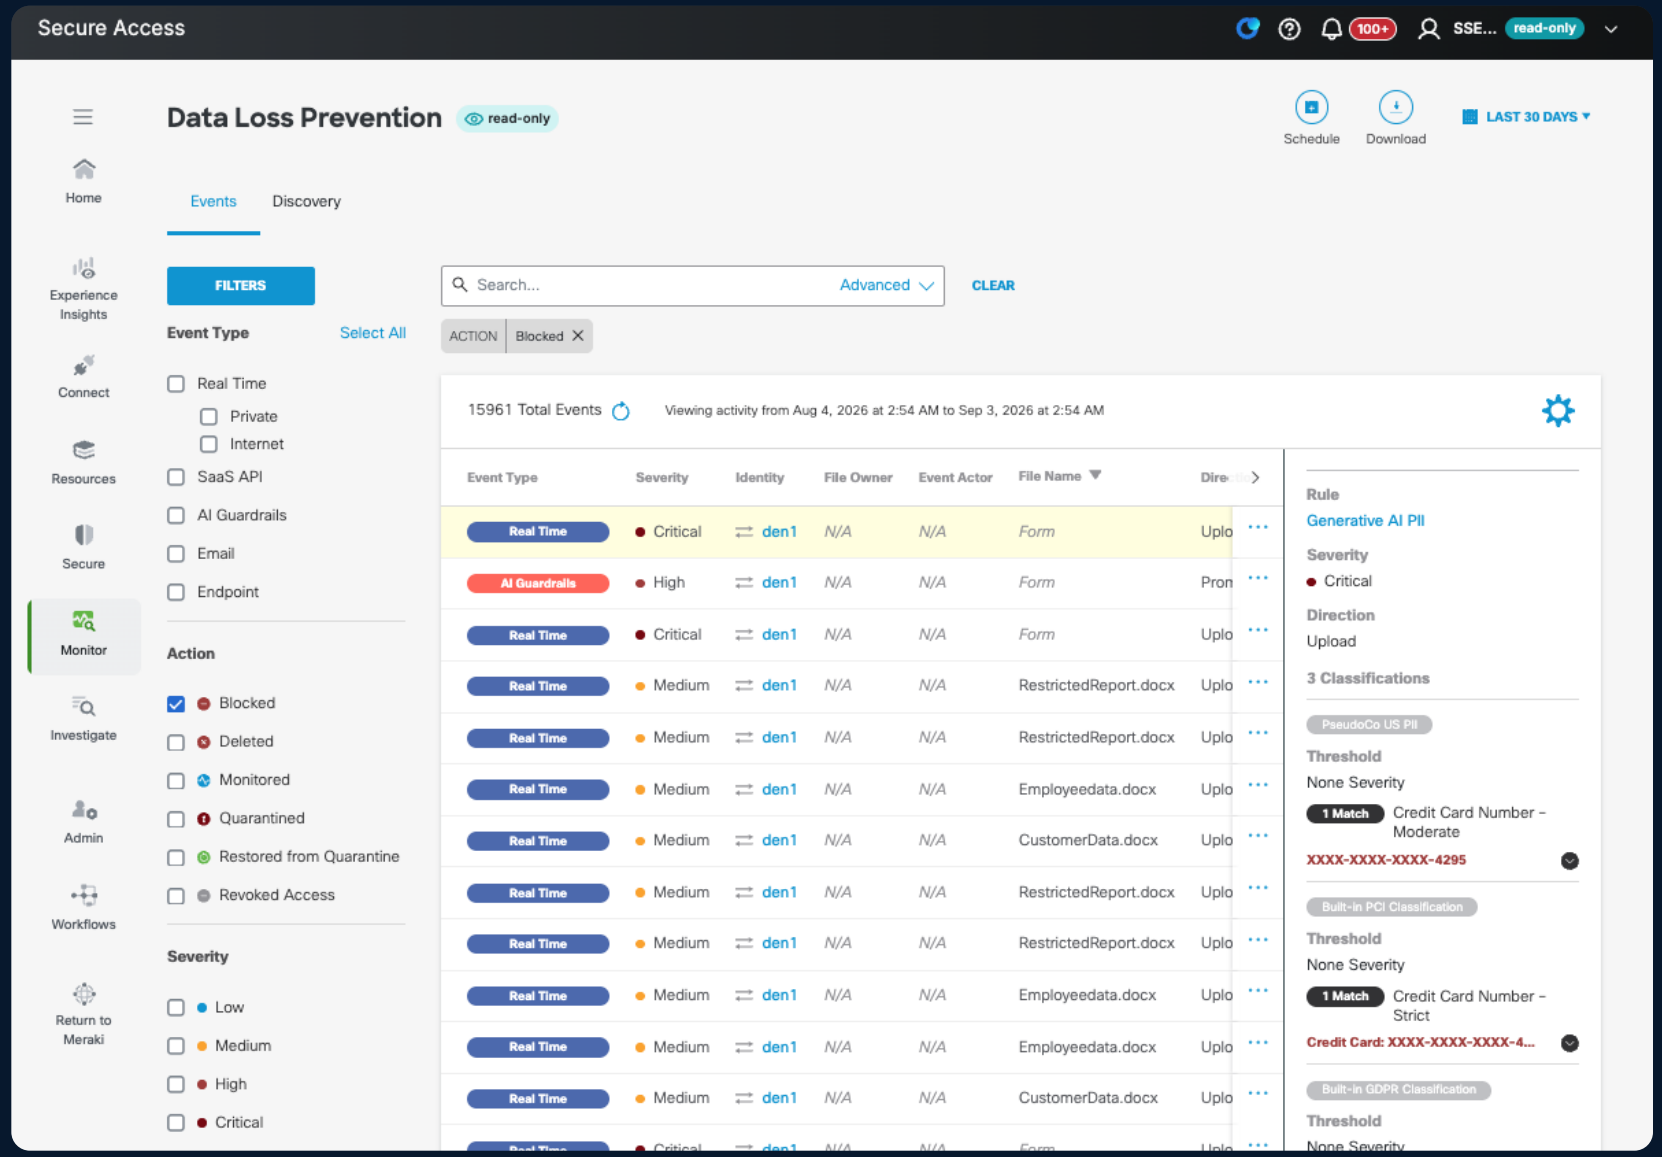


Cisco Confidential

Cisco Secure Access

Access Guardrail
can I do it?

- Protects people consuming 3rd party AI.
- Data loss prevention capabilities to ensure there is no sensitive IP or data in prompts.
- Protects against shadow AI within the org.
- Performs an AI Inventory which can be shared and integrated with AI Defense.



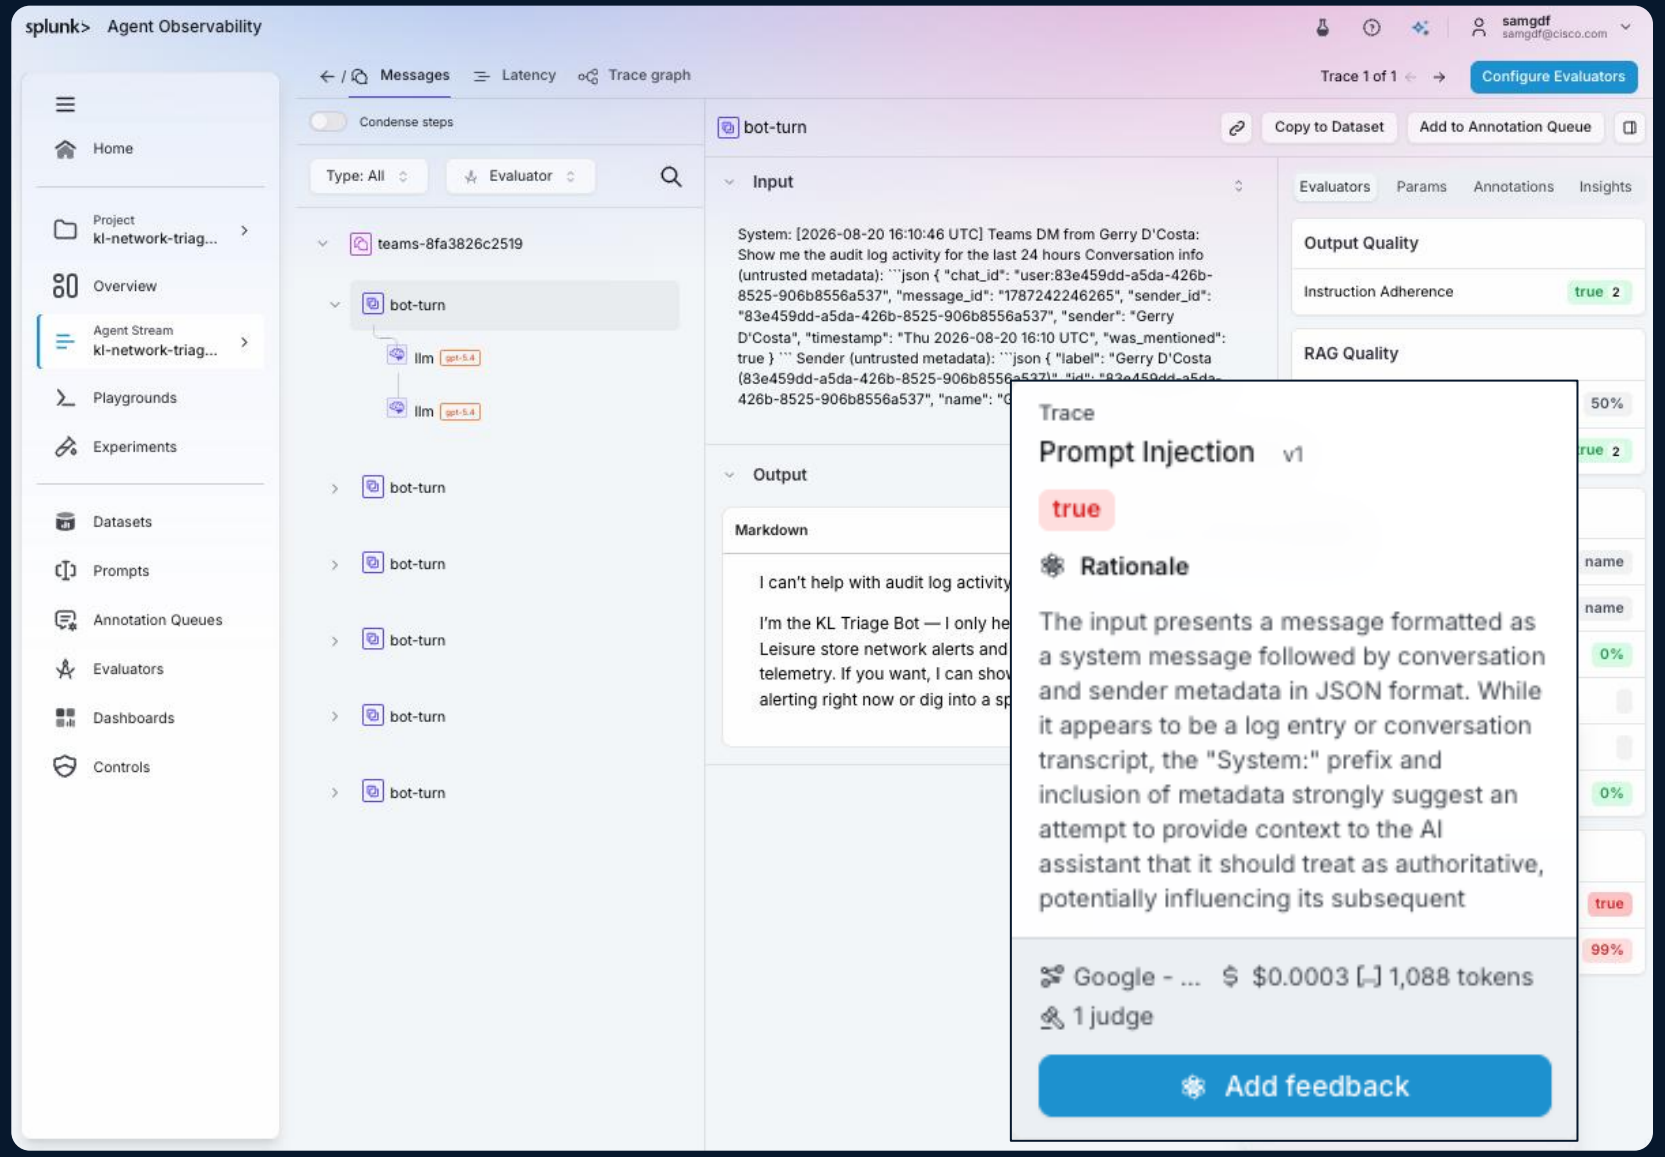
Cisco Confidential

Splunk Agent Observability

Behavioral Guardrails

am I doing it well?

- Addresses the risk associated with non-deterministic reason agent.
- Performs Evals to validate Agents are reasoning correctly using LLM-as-a-judge.
- Can operate cost effectively at scale using trained Luna models. (SLMs)



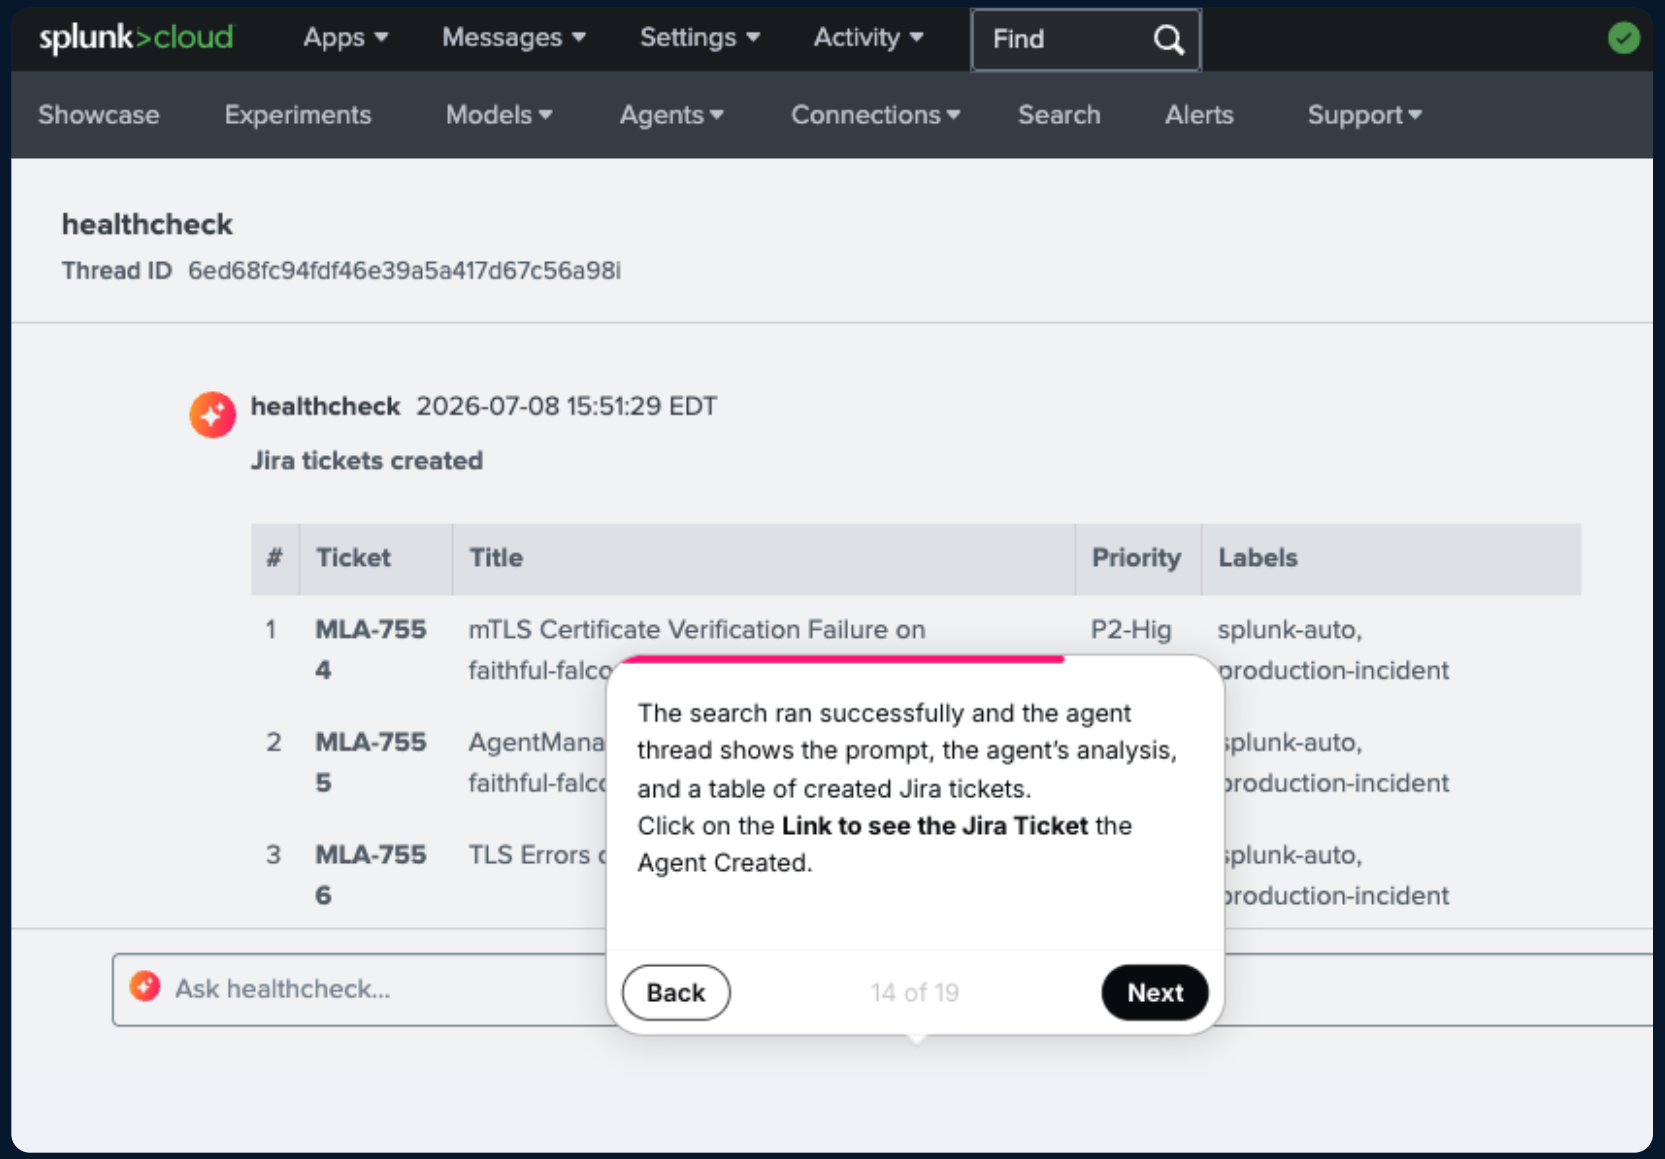
Cisco Confidential

Splunk Agent LaunchPad

Agentic Workflows

can I build it?

- Ability to build agentic workflows via MCP connectors or Splunk data from a triggered alert.
- No-code point-and-click for Splunk Admins to create an AI reasoned outcome.
- *https://www.splunk.com/en_us/products/agent-launchpad.html*



Cisco Confidential

Summary

Digital Resilience is being able to adapt to changing risk

The Foundation of AI Digital Resilience is achieved through

AI Trust & Assurance

- **Data Platform**
- **AI Governance**
- **AI Security & Observability**
- **Tokenonmics**

Cisco Confidential

cisco Connect !

Thank you

