WSO2CON NORTH AMERICA

May 20 - 22, 2026 | Austin, Texas, USA

Who are you: Agent Identity To the Rescue

Ayesha Dissanayaka

Associate Director / Architect

1950

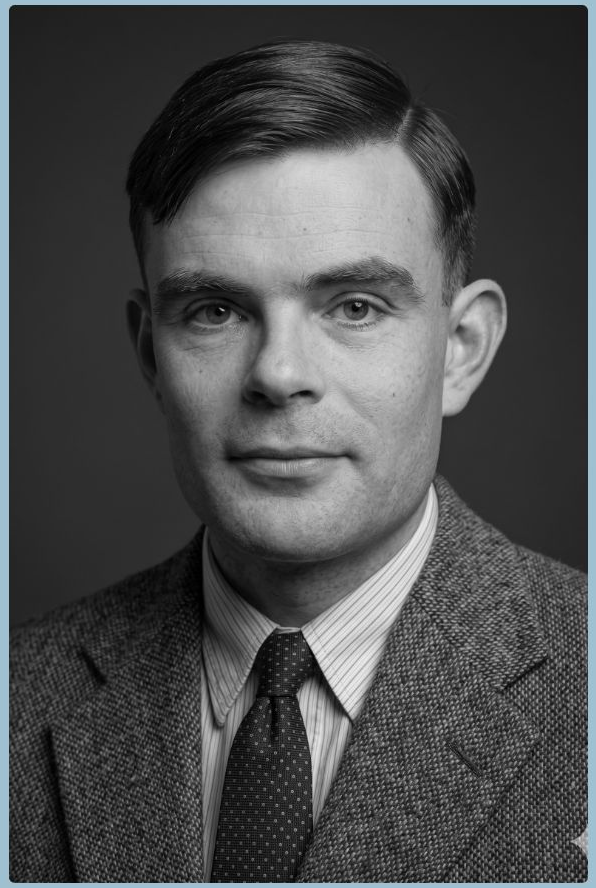

“Can machines *think*?”

The Imitation Game was a contest about a machine pretending to be someone it's not.

COMPUTING MACHINERY & INTELLIGENCE · A. M. TURING, 1950

Can machines think?

~~think~~

Can machines be *trusted*?

TRUST IS NO LONGER A PHILOSOPHY PROBLEM · IT'S AN INFRASTRUCTURE PROBLEM

WSO2CON NORTH AMERICA

May 20 - 22, 2026 | Austin, Texas, USA

Who are you: Agent Identity To the Rescue

Ayesha Dissanayaka

Associate Director / Architect

40%

of enterprise applications will embed task-specific AI agents by the end of 2026 — up from less than 5% in 2025.

GARTNER

Use of AI by respondents' organizations, % of respondents

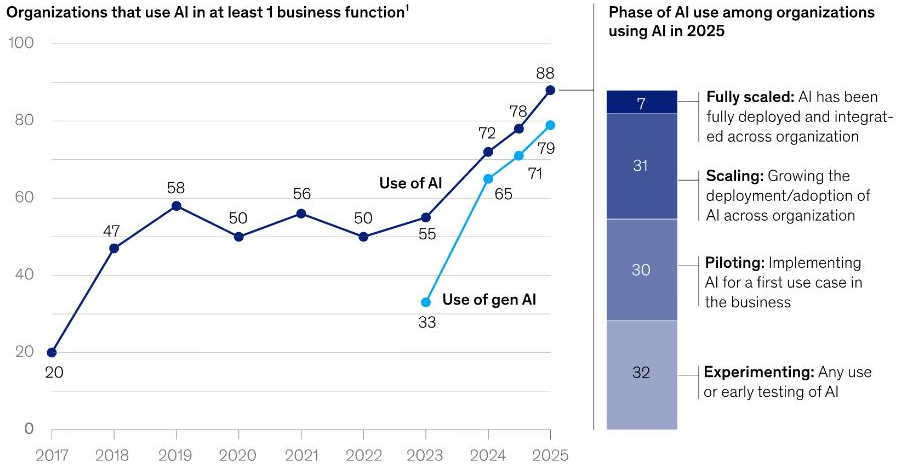

MCKINSEY

97%

AI BREACHES LACKED ACCESS CONTROLS

IBM

80%

EMPLOYEES USING SHADOW AI

UPGUARD

$670K

ADDED BREACH COST IN 2025

IBM

Agents are the *new Taskforce*

They call APIs. They access data. They make decisions. They delegate to other agents. If you wouldn't give a contractor an unscoped, never-expiring key — why are you giving one to them?

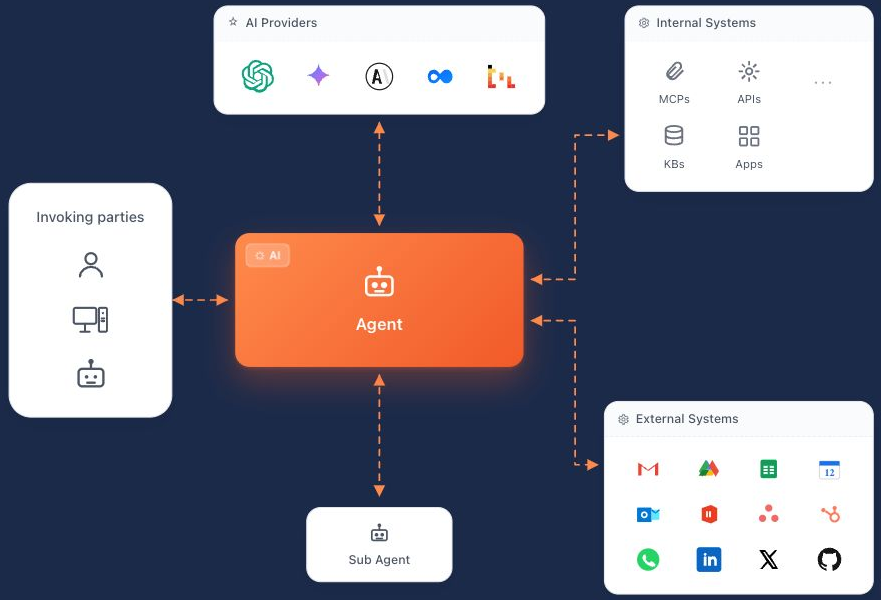

Traditional IAM was built around *humans*. *Agents* are *not* humans.

NO SINGLE CONSTRUCT

Subject, client, and resource — all at once.

NO SCRIPT

Goal-oriented, not predefined execution.

CLASS VS. INSTANCE

Spans models, runtimes, vendors

SPEED OF CHANGE

Tomorrow's tool catalog isn't on today's list.

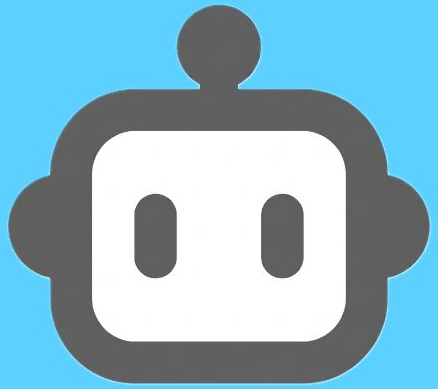

Meet *Agent X.*

An ordinary enterprise AI agent. The one your team spun up last quarter.

Q1.	*"Who registered you?"*
Q2.	*"How do we know it's you?"*
Q3.	*"What are you allowed to do?"*
Q4.	*"Can you prove what you did?"*

The *4A's* framework for *Agent IAM*.

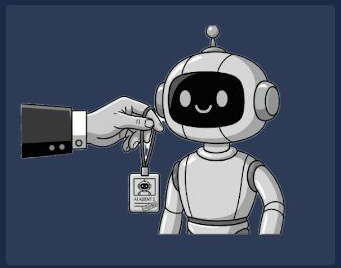

*A*dminister

"Who registered you?"

- Registration
- Ownership
- Lifecycle
- Deprovisioning

*A*uthenticate

"How do we know it's you?"

- Attestation
- mTLS & workload IDs
- Short-lived tokens
- Zero-trust

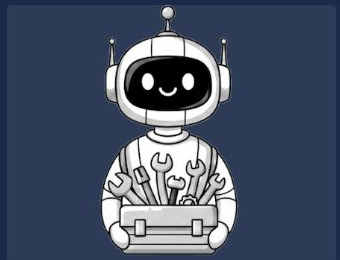

*A*uthorize

"What are you allowed to do?"

- Least privilege
- Just-in-time
- Tool-level grain
- Human-in-the-loop

*A*udit

"Can you prove what you did?"

- Action trails
- Decision provenance
- Compliance evidence
- Anomaly detection

Every *edge* is a place to *Authenticate, Authorize & Audit*.

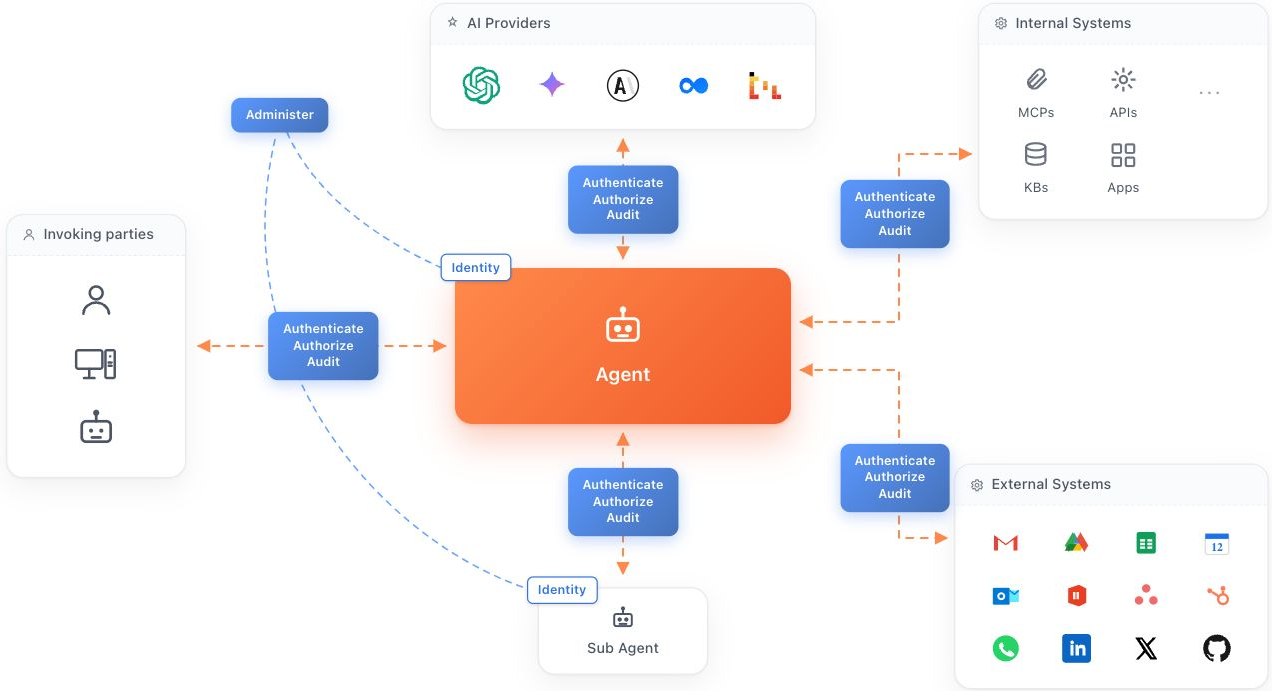

From inventory to *autonomy*.

PHASE 01

Crawl

Know what you have before you govern it.

- Inventory every agent identity
- Classify by blast radius
- Name a human owner per agent
- Educate developers and teams

PHASE 02

Walk

Centralize the basics. One front door.

- Unified Agent Identity
- Streamline Access
- Short-lived tokens, auto-rotation
- Baseline audit on every call

PHASE 03

Run

Dynamic. Contextual. Complete.

- Just-in-time authorization per task
- Human-in-the-loop on sensitive actions
- Provenance-grade audit trail
- Egress gateway with policy enforcement

Do

Short-lived tokens.
Minutes, with automatic rotation.

Scoped credentials per task.
Mint at need. Expire with the work.

Human owner per agent.
A name on the credential.

Audit-first thinking.
If you can't query it, it didn't happen.

Don't

Hardcoded secrets.
Long-lived keys in env vars — a breach with a calendar invite.

Over-permissioning.
"Just give it admin." The most common, costliest mistake.

Skipping audit.
"We'll add it later" → "we don't know what happened."

Static RBAC.
Roles on a system whose scopes change every prompt.

OWASP Agentic *Top 10* at a glance.

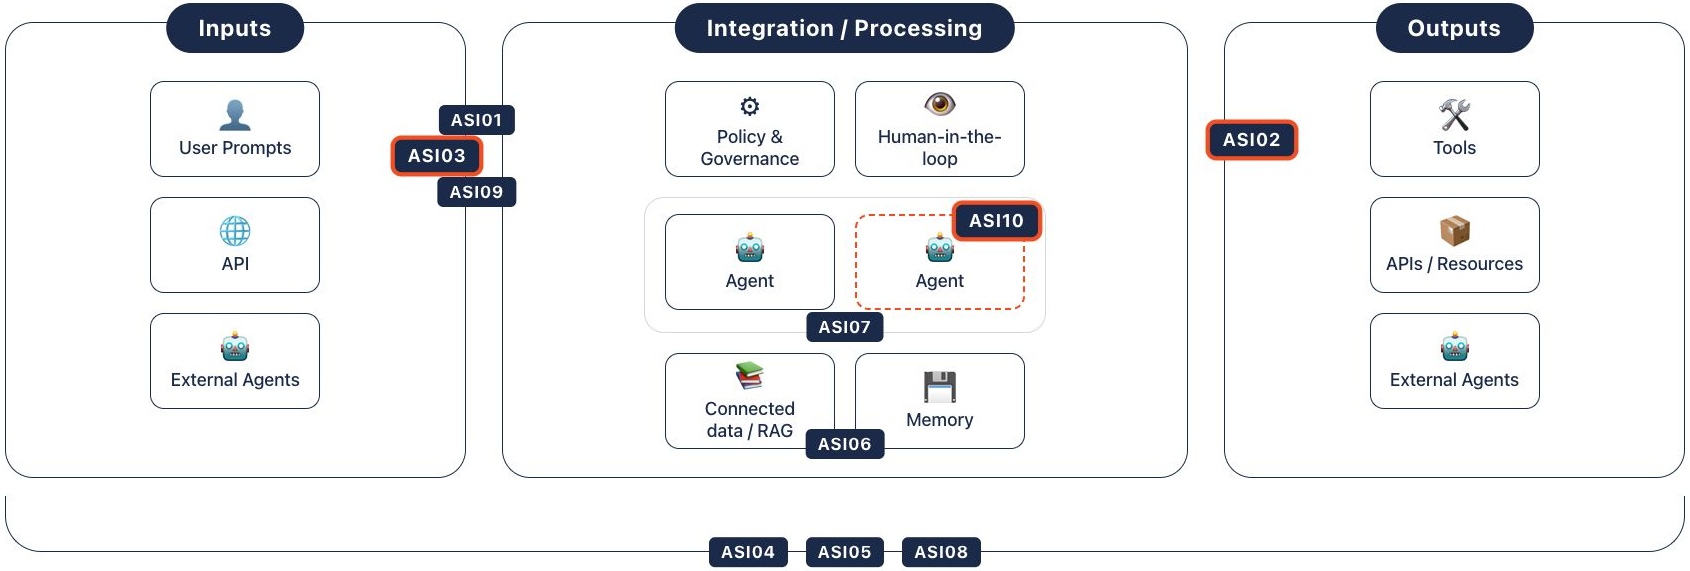

ASI01: Agent Goal Hijack

ASI02: Tool Misuse & Exploitation

ASI03: Identity & Privilege Abuse

ASI04: Agentic Supply Chain Vulnerabilities

ASI05: Unexpected Code Execution (RCE)

ASI06: Memory & Context Poisoning

ASI07: Insecure Inter-Agent Communication

ASI08: Cascading Failures

ASI09: Human-Agent Trust Exploitation

ASI10: Rogue Agents

What IAM *cannot* solve.

i. **Prompt injection**

Auth was clean. The instruction was the attack.

ii. **Hallucination**

A confidently wrong answer is not an access problem.

iii. **Training data poisoning**

Compromise months before the agent authenticates.

iv. **Adversarial inputs**

The credential is valid. The behavior is not.

Standards across the agentic *stack*.

MCP *A2A* A2P ACP

OAuth 2.1 OIDC *Token Exchange* CIBA

SPIFFE RAR

AuthZEN OPA Cedar XACML

NIST AI RMF AAIF · Linux Foundation EU AI Act

Emerging IAM *practices*.

01

Agent federations

Trust frameworks across enterprises.

02

Trust scoring

Risk-weighted authz on behavioral history.

03

Capability-based access

Right to act, not right to be — composable.

04

AI-governed IAM

Models watching other models for drift.

05

Intent-based authz

Authorize the intent, not just the call.

06

Cross-protocol bridges

SPIFFE ↔ OAuth ↔ MCP.

Your Back-to-Office Checklist.

01	Agents need *first-class identities*, not shared API keys.
02	Use the *4A lens* to build your Agent Identity solution.
03	Protocols are evolving — *build for adaptability*.
04	IAM is *necessary, not sufficient*.
05	Start now. *Crawl, walk, run*. But do it fast.

WSO2CON NORTH AMERICA

May 20 - 22, 2026 | Austin, Texas, USA

Thank You!

Ayesha Dissanayaka

Associate Director / Architect

Resources & References

wso2.com/identity-platform/agent-id/