

September 29, 2026



Why Agent Control Planes Matter

(and what one actually looks like)



Tori Jeffcoat

Director, Product Marketing - AI



Evaluation Summary

Average Score

72.97%

342 total evaluations

Trace 26 traces
4 evaluators · 16.3% skipped

Agent 26 agent executions
3 evaluators

LLM 40 LLM invocations
4 evaluators

Provider Template



Anthropic



AWS



Gemini



Microsoft

Set up

Select Build

hotel-booking-agent-build-1bddb9dd

ac71eb2f 06 Feb 2026



Configure & Deploy

Icebreaker:
What is the *best* ice
cream flavor and
topping?



WSO2 Agentic Enterprise Fabric

Trusted AI Governance across the stack



Govern every **agent**



Govern every **API & AI service**



Govern every **connection**



Govern every **identity**



Govern every **workload**



Today's Agenda

What is an agent control plane?

Key requirements for agent control

WSO2 Agent Manager

Q&A





What is an agent control plane?



Rania Khalaf

Chief AI Officer

What makes agents **powerful**

— *the combination of non-determinism and autonomous actions* —

is what makes them **hard to govern**



Without dedicated agent controls, agent programs often end up in one of two places

STUCK IN PILOT

Teams can't get agents into production

40%

of agentic AI projects
predicted to be cancelled
by EOY 2027

Gartner

AGENTS OUTPACE CONTROLS



Leaders can't get a handle on agent sprawl

#1

Concern facing CIO's is
Agent Sprawl

Gartner





55.4% increase in **documented** AI incidents between 2024 (233) and 2025 (362)

"The AI Index 2026 Annual Report," AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Stanford, CA, April 2026



Governance must be redesigned for probabilistic systems

Agent Action

1.Decides.

The agent is told what to achieve, not what to do.

2.Acts.

The agent picks tools, arguments, and timing on its own.

3.Not a service, not a user

A third actor with its own intent.

4.Non-deterministic.

Same prompt, different completion.

Governance Consequence

Governance cannot enumerate the agent's actions in advance.

The action surface is combinatorial, not a list.

Identity, Scope, Audit per agent. On-behalf-of chain.

Behaviour is a distribution, not a contract.



**You cannot run a
probabilistic agent on
infrastructure built
for deterministic
code.**



Spiderman, © 2002 Columbia Pictures Industries, Inc. and Marvel Characters, Inc. All Rights Reserved



Agent Control Category

Agent Control Plane

FORRESTER®

AI Agent Management
Platform (AMP)

Gartner®

Multi-agent estate

FinOps


Security

Visibility

Identity



Governance across every level of adoption



Agent Control Plane: the ice cream, not the ingredients



Key requirements for agent control

Agent Management Solution: Key Capabilities

Federated agent management

Any framework



CrewAI



LangChain



AWS Strands



Ballerina



Microsoft Agent Framework

Any model



Anthropic



Gemini



AWS Bedrock



Mistral



OpenAI



Azure AI Foundry



Your Own Model

Any deployment



Agent Manager



AI Agents



External



AI Agents




AI Agents



Agent Management Solution: Key Capabilities

Full Lifecycle Management



Agent Management Solution: Key Capabilities

Agent Identity & Security

**Customer Support Agent** PLATFORM-HOSTED

 Edit Agent

 Framework: CrewAI

 1 day ago

 Admin User

Roles

ROLE

 **Guest Services**
Operations MCP: No access

Groups

GROUP

 **Guest Facing Agents**

Credentials

Authentication details for this agent to securely access applications and API resources.

AGENT ID

.....-.....-.....-.....-.....





Agent Management Solution: Key Capabilities

Governance & Guardrails

LLM and MCP controls

[← Back to LLM Providers](#)

 **Hotel Concierge Guarded Provider** OpenAI v1.0
Guardrailed OpenAI provider for Hotel Concierge Agent runtime traffic

[Overview](#) [Connection](#) [Deployment](#) [Access Control](#) [Security](#) [API Keys](#) [Rate Limiting](#)

Global Guardrails
Applies for all resources

[Prompt Decorator \(v1.0.2\)](#) [+ Add Guardrail](#)

 **Hotel Operations MCP** v1.0.0
Internal ops tools for the Grand Meridian hotel: room-status updates and maintenance tickets.
🕒 about 1 month ago
🕒 22 days ago

[Overview](#) [Connection](#) [Manage Tools](#) [Security](#) [Rewrite](#) [Policies](#)

[Tools](#) [Scopes](#)

Tool	Scopes
update_room_status	manage  Assign scopes...
create_maintenance_ticket	manage  Assign scopes...

Agent level controls

 **Hotel Concierge Guarded Provider** | OpenAI 

Deployed 29 days ago
Rate Limiting: Not configured
Guardrails: Prompt Decorator

Agent Level

Guardrails [+ Add Guardrail](#)
Add safety policies to enforce consistent protections.

[Prompt Decorator](#) 

[PII Masking Regex](#) 

Provider Guardrails
Configured on the LLM provider itself, applied in addition to the guardrails above.

[Prompt Decorator](#)

Org level



Agent Management Solution: Key Capabilities

Monitoring & Evaluation

Agent Performance



Evaluation Summary

Average Score

72.97%

342 total evaluations

Trace 26 traces
4 evaluators - 16.3% skipped

Agent 26 agent executions
3 evaluators

LLM 40 LLM invocations
4 evaluators

Run Summary

[View all runs](#)

Window Start	Run Score
03/03/2026, 18:12:24	72.97%
03/03/2026, 18:12:24	72.05%



1.0 Now GA

WSO2 Agent Manager

The open enterprise control plane to securely deploy, manage, and govern AI agents at scale.

Self-hosted or SaaS

Built for Humans & Agents



LANDSCAPE REPORT

The Agentic Control Plane Solutions Landscape, Q2 2026

Forrester's Overview Of 33 Vendors

June 24, 2026

By Leslie Joseph with Frederic Giron, Laura Demarki, Bill Nagel

FORRESTER

Summary

You can use agentic control plane solutions to govern a fragmented agent estate, contain the risk that agent autonomy introduces, and keep production agents observable and accountable. But to realize these benefits, you'll first have to select from a diverse set of vendors that vary by size, type of offering, geography, and use case differentiation. Enterprise architecture and AI leaders should use this report to understand the value they can expect from an agentic control plane solution vendor, learn how vendors differ, and investigate options based on size and market focus.



How WSO2 Agent Manager Works





Live Demo

WSO2 Agent Manager



Asanka Abeyweera

Associate Director & Architect



Closing



Rania Khalaf

Chief AI Officer

Go faster

Major shifts happen seemingly overnight

 TypeSafe AI

Company News

Sep 15, 2026

Introducing System One Models & Jev

 Fortune

Uber burned through its entire 2026 AI budget in four months. Now its COO is questioning whether it's worth it

The rideshare giant's COO says "it's very hard to draw a line" between rising AI costs and useful features for customers.

26 May 2026



 The GitHub Blog

GitHub Copilot is moving to usage-based billing

Starting June 1, your Copilot usage will consume GitHub AI Credits.

27 Apr 2026



Pump the brakes

Resilience and regulatory concerns can stop you in your tracks

AI



Announcements

Statement on the US government directive to suspend access to Fable 5 and Mythos 5

Jun 12, 2026

The US government, citing national security authorities, has issued an export control directive to suspend all access to Fable 5 and Mythos 5 by any foreign national, whether inside or outside the United States, including foreign national Anthropic employees. The net effect of this order is that we must abruptly disable Fable 5 and Mythos 5 for **all** our customers to ensure compliance. **Access to all other Anthropic models will not be affected.**

We received the directive from the government today at 5:21pm (ET). The letter did not provide specific details of its national security concern. Our understanding is that the government believes it has become aware of a method of bypassing, or “jailbreaking” Fable 5. We reviewed a demonstration of this specific technique being used to identify a small number of previously known, minor vulnerabilities. These vulnerabilities all appear relatively simple, and we have found that other publicly-available models are able to discover them as well without requiring a bypass.

Anthropic's posture with respect to Fable's safeguards, as laid out in our launch [blog post](#), is the following:

Forbes

INNOVATION > CYBERSECURITY

Microsoft Can't Keep EU Data Safe From US Authorities

By [Emma Woollacott](#), Senior Contributor. © Emma Woollacott is a freelance jo... ▾

[Follow Author](#)

Published Jul 22, 2025, 07:56am EDT

In sworn [testimony](#) before a French Senate inquiry into the role of public procurement in promoting digital sovereignty, Anton Carniaux, Microsoft France's director of public and legal affairs, was asked whether he could guarantee that French citizen data would never be transmitted to U.S. authorities without explicit French authorization. And, he replied, "No, I cannot guarantee it."

He said that the company resisted requests from the US authorities "when they are not well-founded", but that under the U.S. Cloud Act, U.S. companies can be forced to hand over data, regardless of where it is stored.




Creating a need to architect for resilience, flexibility, and control

Score Breakdown by Model

Model	Tone	Count
openai/gpt-4o	69.7%	19
anthropic/sonnet-5	74.3%	23

Agent resiliency: Switch models without rebuilding governance



Control Plane and Data Plane Separation
Allows for centralized governance with distributed execution



Thank you! Questions?

Trusted AI Governance with WSO2

Open • Resilient • Yours

Learn more
about WSO2
Agent Manager

