affinidi

2026 MARKET TREND REPORT

The Cross-Boundary Gap in the Agentic Economy

Why trust, identity, and accountability now have to travel with every AI agent that acts beyond its boundaries

The harder problem in 2026 is that AI agents need to step beyond their own 'walls' to fulfil their full potential and build trust across clouds, organisations, platforms, and environments. Today, they can't, because the tools that govern them stop at the organisational boundary. As they cross into those environments, identity, trust and accountability have to travel with them.

July 2026

AUTHOR

Brian Kim, VP of Strategy & Growth, Affinidi

With independent commentary from CardInfoLink

Contents

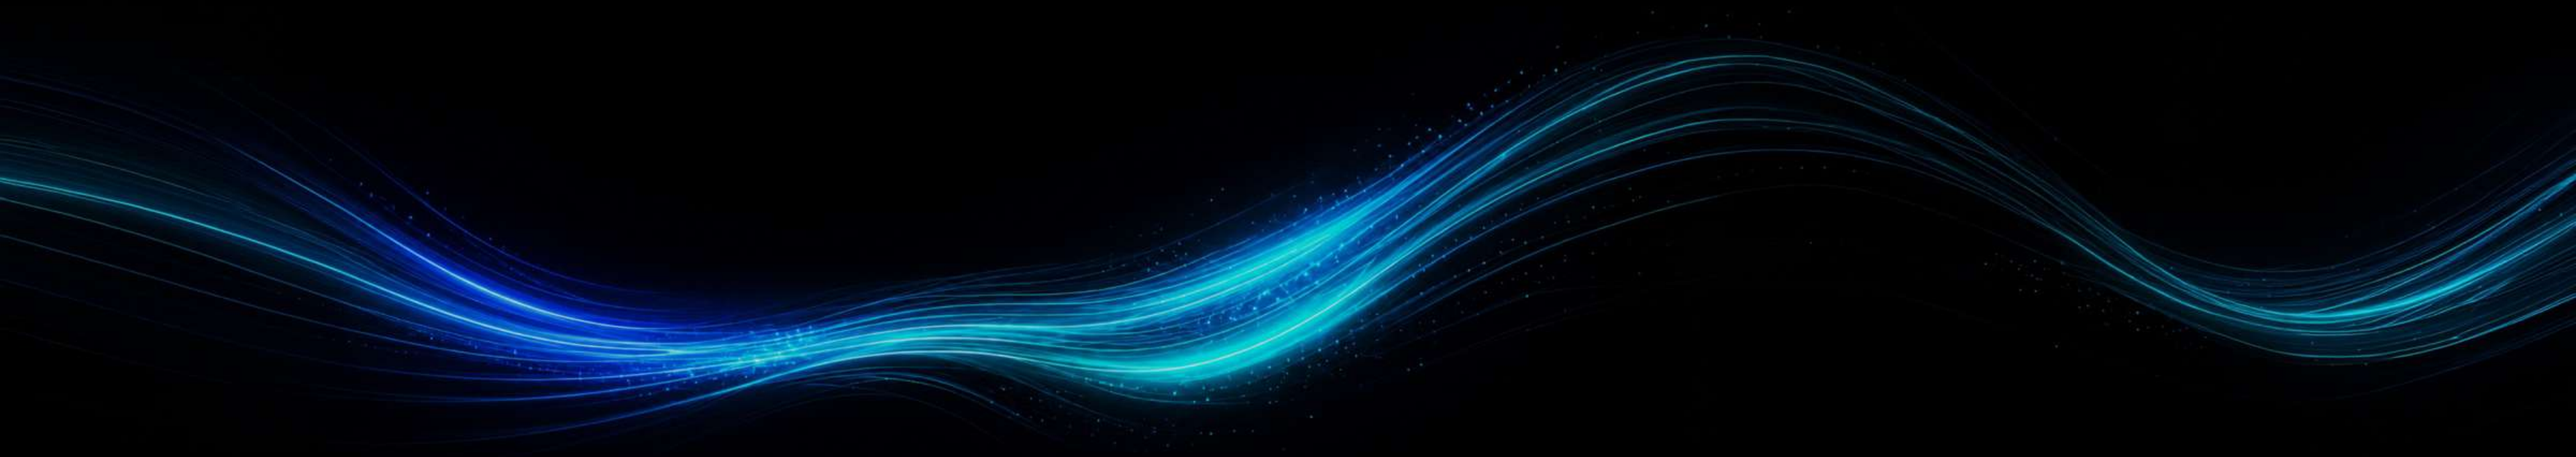

01

Foreword: AI agents are ‘leaving the building’

For most of the past year, an AI agent was something that operated within one company. It read that company’s data, used that company’s tools and answered to that company’s staff. Governing it was an internal exercise, and the market was built accordingly.

That assumption is broken.

The same AI agents are now being extended to act outside the environment that created them, collaborating with other tools or agents, negotiating with a supplier’s AI agent, buying from a merchant, and settling a payment with a counterparty they do not control.

The moment an AI agent crosses that line, the internal control model runs out of road, and the AI agent is now in a territory where its home platform has no visibility, and the other side cannot prove it is safe to trust.

The proof is in the plumbing

The clearest evidence that AI agents are crossing boundaries is that academia, consortia, and industry players have spent the past year building the rails for exactly that. Each of these efforts had to invent a way for an AI agent to prove who it is and who stands behind it to a party that does not share its platform. In other words, portable trust is being built out of necessity.

While various efforts exist to solve this problem, they also share a common limitation. Each provides a way to move and connect AI agents, but not the ability to prove who the agent is or who is accountable for it outside of its 'walls'. That identity and accountability layer outside of its boundaries is the open question this report examines.

NHI (Non-Human Identities) already outnumber humans by over 100:1 (Palo Alto, 2026), with an estimated ~70% of NHIs being AI agents. And unsurprisingly, the use cases are also increasingly cross-boundary.

CROSS-BOUNDARY EFFORT	SINCE	WHAT IT IS	ADOPTION SIGNAL (AS OF JUN 2026)
Linux Foundation A2A	Jun 2025	Vendor-neutral agent-to-agent interoperability. Launched by Google in April 2025 with 50+ partners, moved to the Linux Foundation in June 2025 for neutral governance	150+ supporting organisations; v1.0 adds signed agent cards
Google AP2	Sep 2025	Open Agent Payments Protocol for agent-to-merchant transactions	60+ partners at launch (Sep 2025), passed 100 by late Oct 2025; donated to the FIDO Alliance (Apr 2026) to stay platform-agnostic
NANDA	Apr 2025	Agent infra protocol / framework by MIT for discovery, identity, trust, verification for AI agents operating across organisations	Largely still experimental with small open-source community, with limited enterprise adoption (Possibly some closed pilots)
SPIFFE	~2016	Open standard for workload (incl. agents) identities across boundaries	Mature production grade protocol and integrated or supported by major hyperscalers and tech vendors; >400 enterprise / tech org contributors

Grades: A2A is verified (Linux Foundation primary). AP2 is vendor announcements dated and corroborated by independent reporting. The 109:1 ratio is based on a vendor survey and varies widely by source and environment (other vendors report ratios ranging from 45:1 to 144:1).

The cross-boundary trust problem is already here. But the agent governance model is still solving last year's problem, stopping at the edge of one tenant.

The question every leader should now ask of any AI agent goes beyond "can we see it?" to "can it prove who it says it is across boundaries, for us to trust it?"

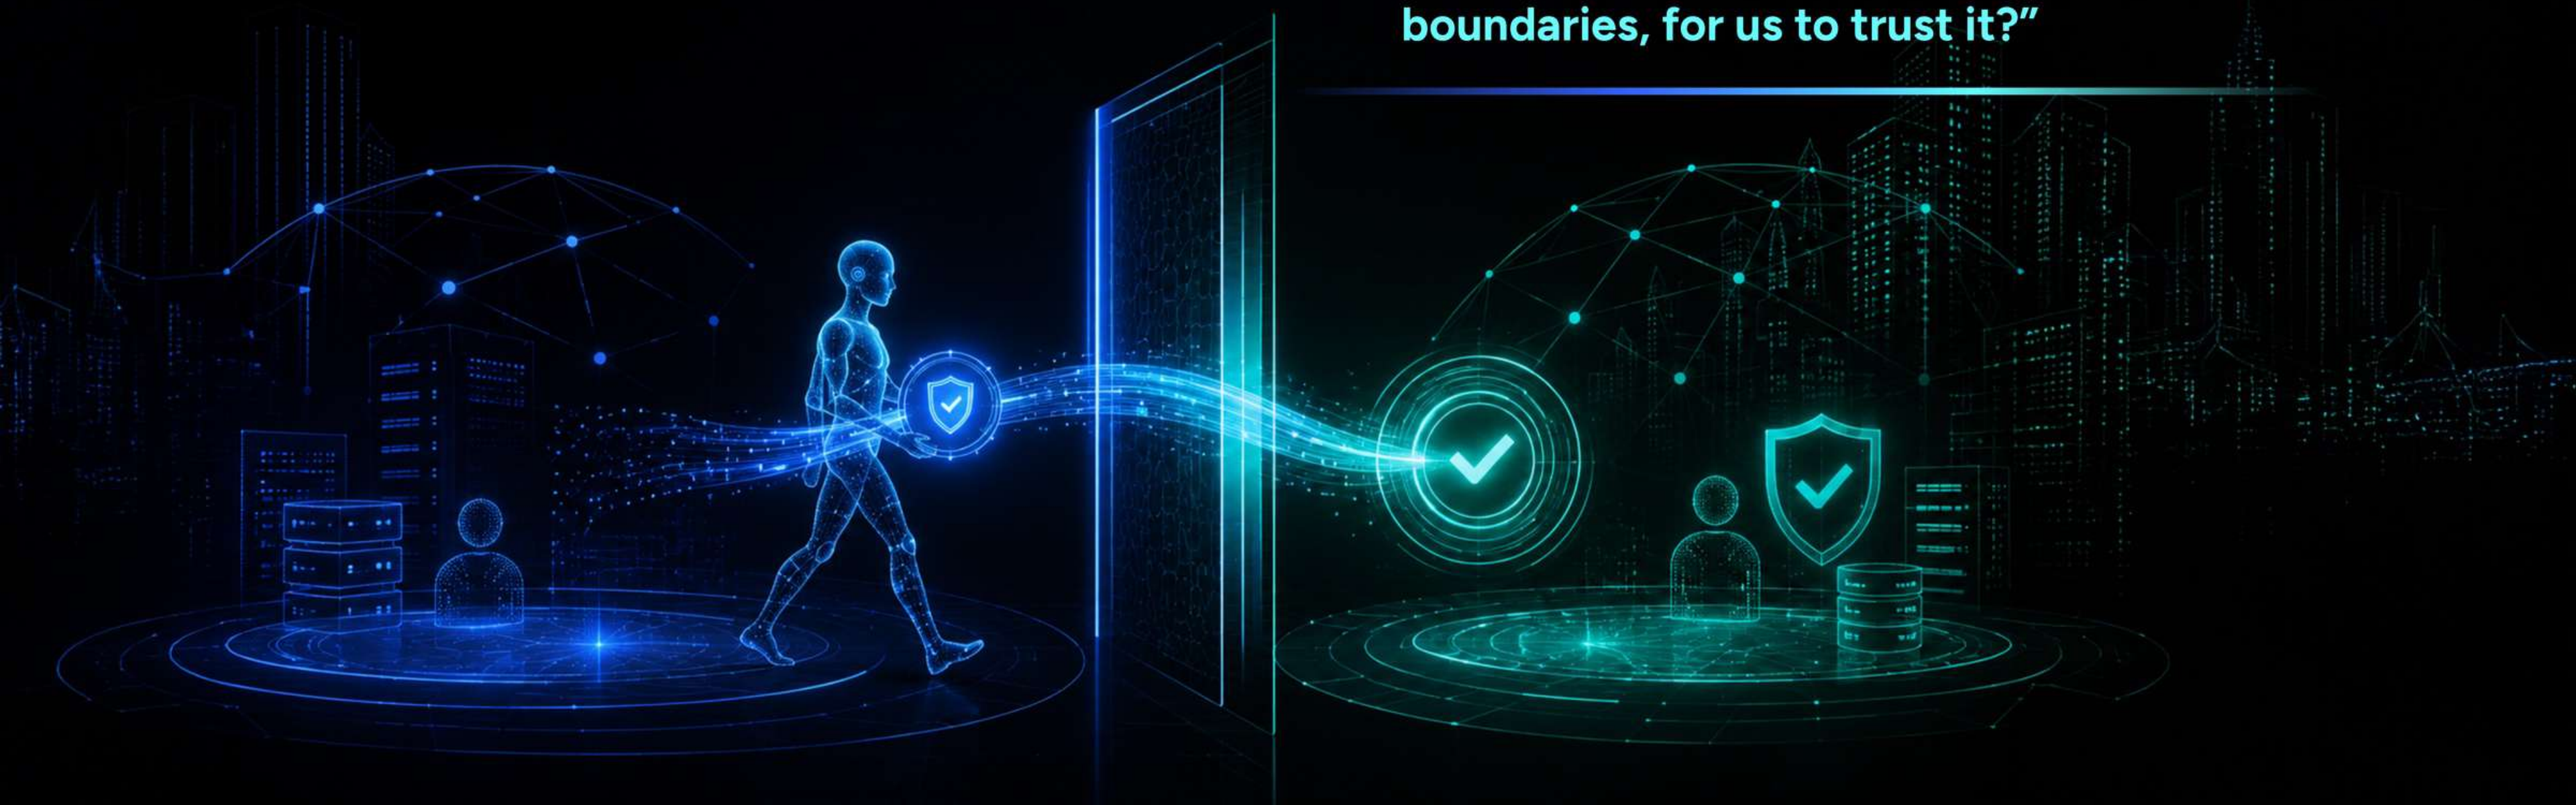

Executive summary

THE SHIFT

Situation

AI Agents are moving into production within enterprises, and the market has quickly responded with solutions to govern and trust agents for internal use cases.

Complication

However, the current offerings are mostly only good for governing agents within their native platforms. Meanwhile, these same AI agents are now being brought to act across organisations, clouds, platforms, and system boundaries, transacting with other third-party agents, services and systems. Most control planes on the market stop at the edge of their native environment.

Question

In complex cross-boundary agent workflows, how do you enable AI agents to collaborate with external actors at scale, with agility, high reliability, and control?

Accountability has to travel with the AI agent. It has to bind to a human principal and remain independently verifiable even after the AI agent leaves the environment that created it.

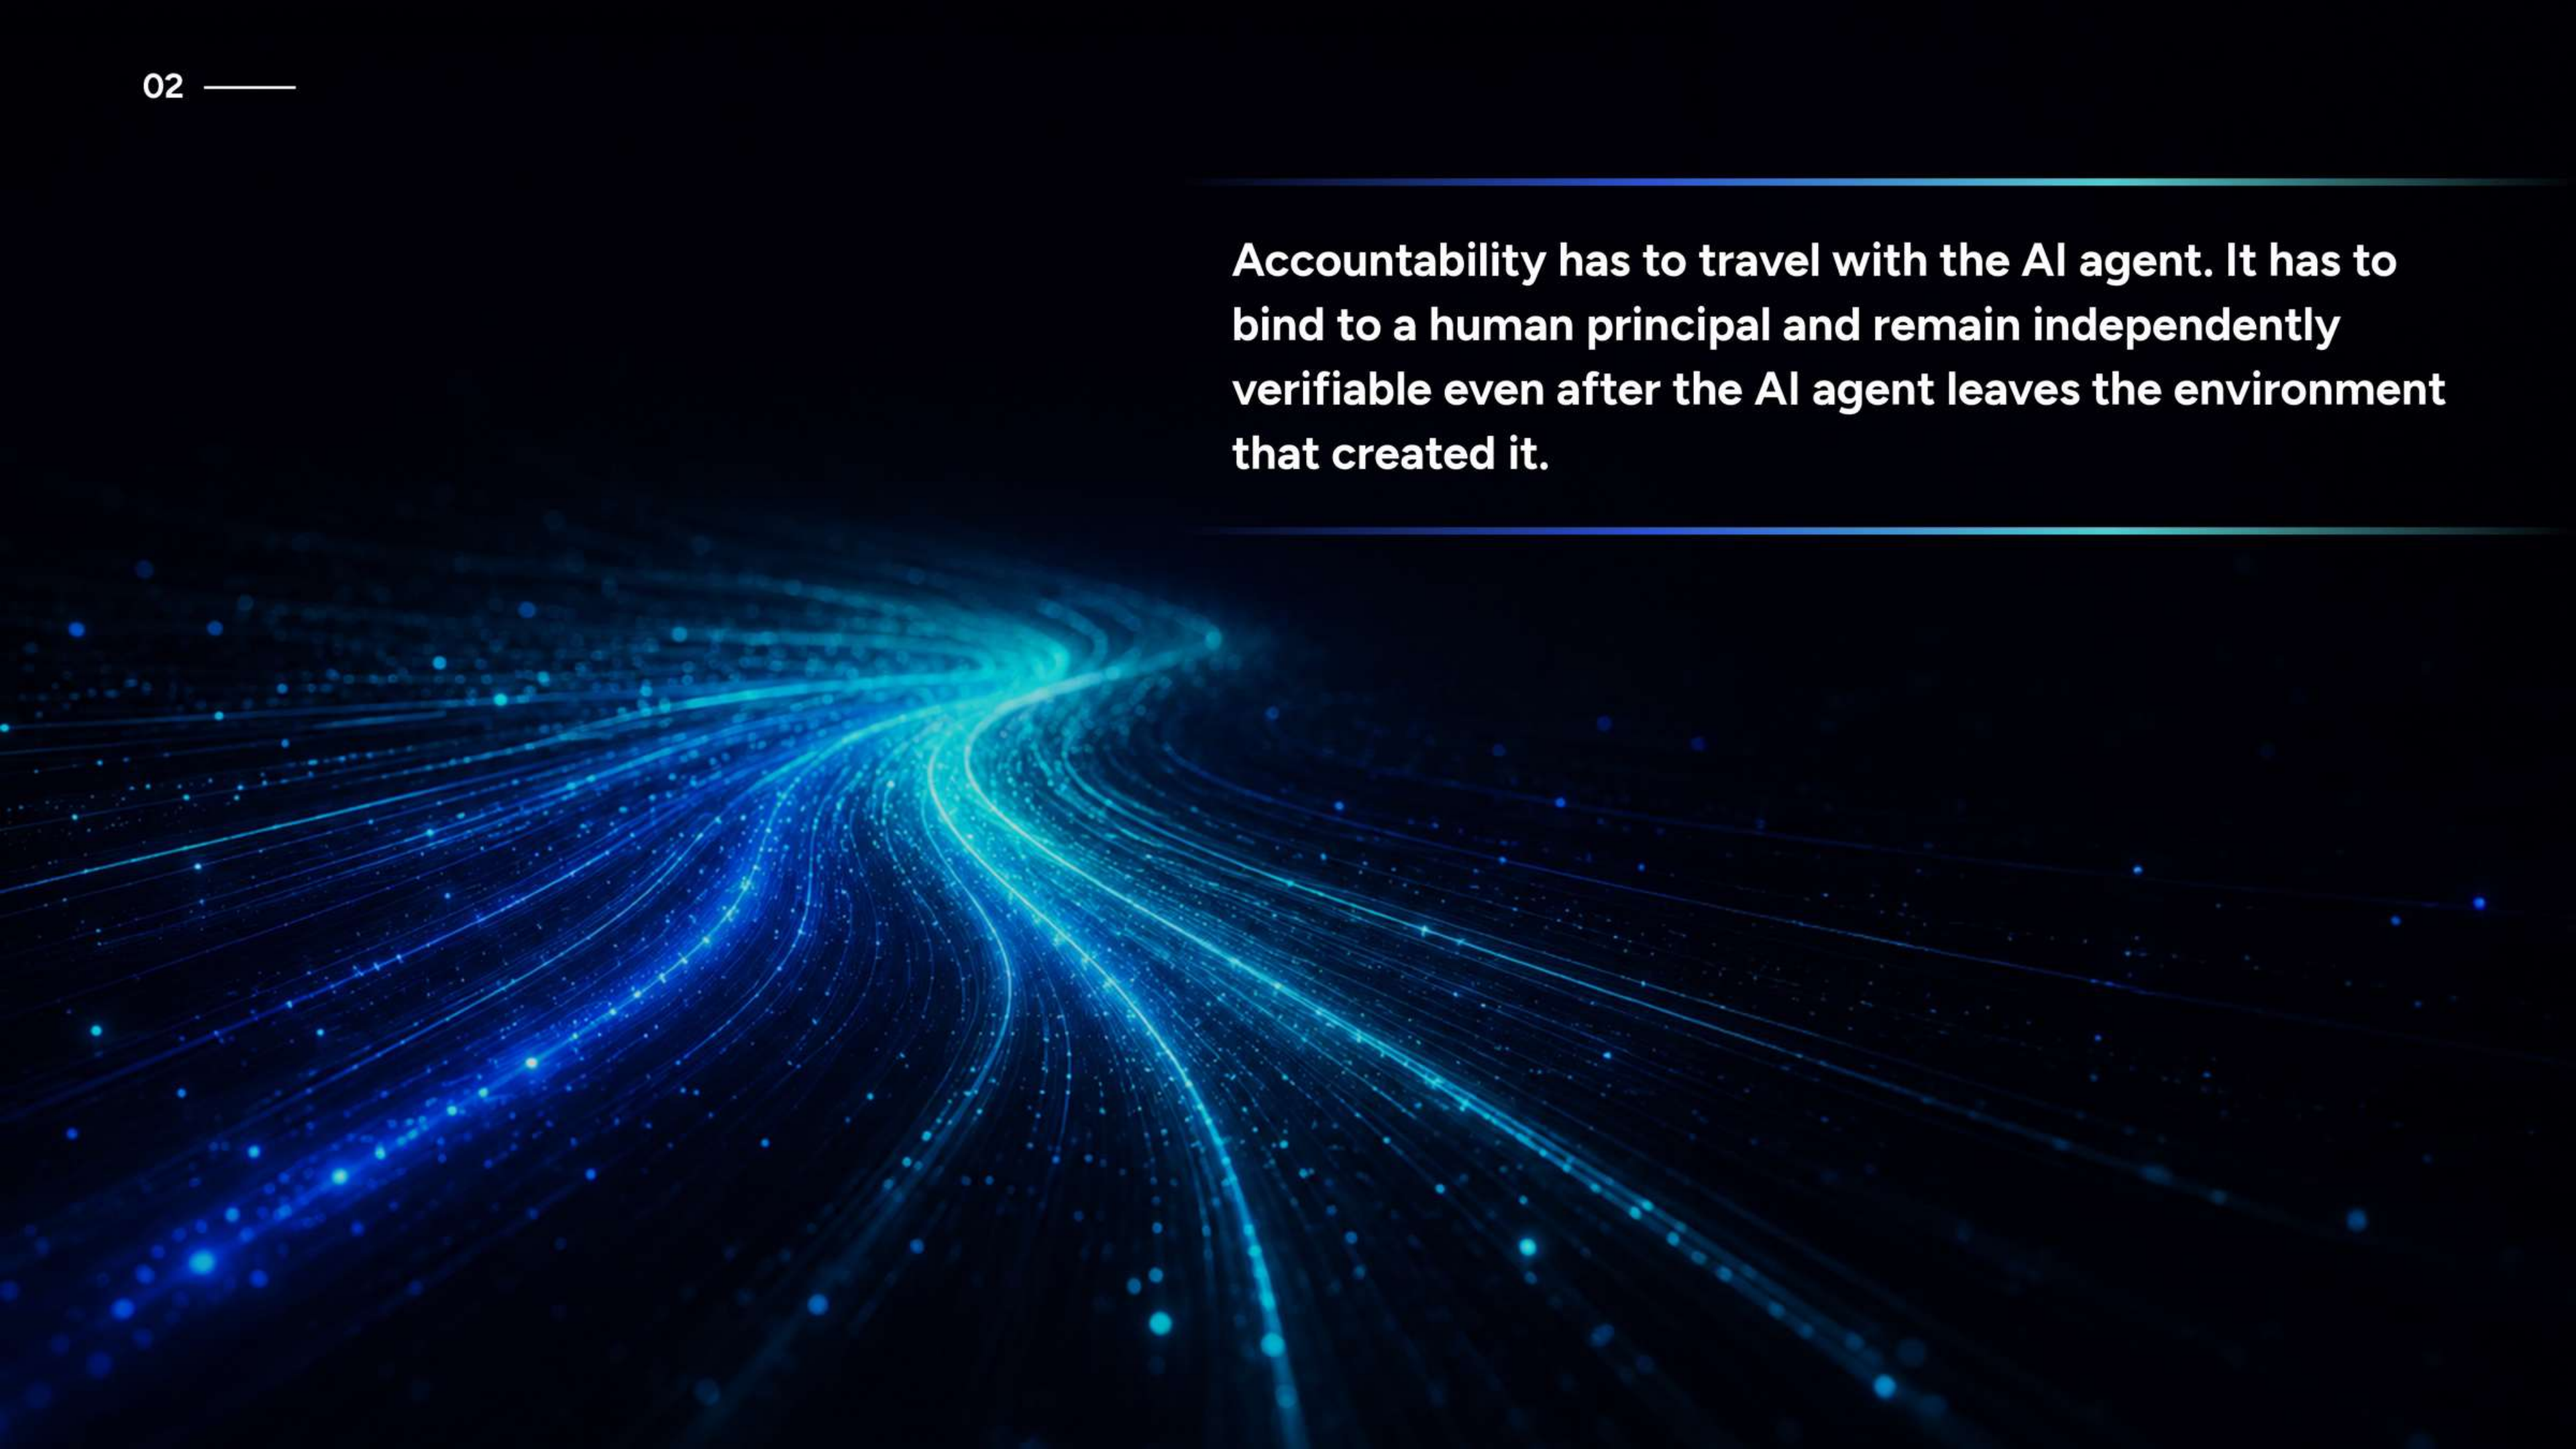

What the evidence shows

1 Adoption is real and so is the failure rate.

Gartner predicts that over **40%** of agentic AI projects will be cancelled by the end of 2027, citing cost, unclear value and inadequate risk controls. The projects most exposed are those that cannot operate beyond a single vendor's walls.

2 AI Agents are already crossing boundaries

The Linux Foundation's Agent2Agent project passed **150** supporting organisations in April 2026.

Google's Agent Payments Protocol has reached 100 committed partners, a measure of industry endorsement rather than live transaction volume, and has been donated to the FIDO Alliance.

Mastercard has shipped agent-commerce rails and run live transactions, and Visa has published its agent-commerce protocol and run pilot transactions.

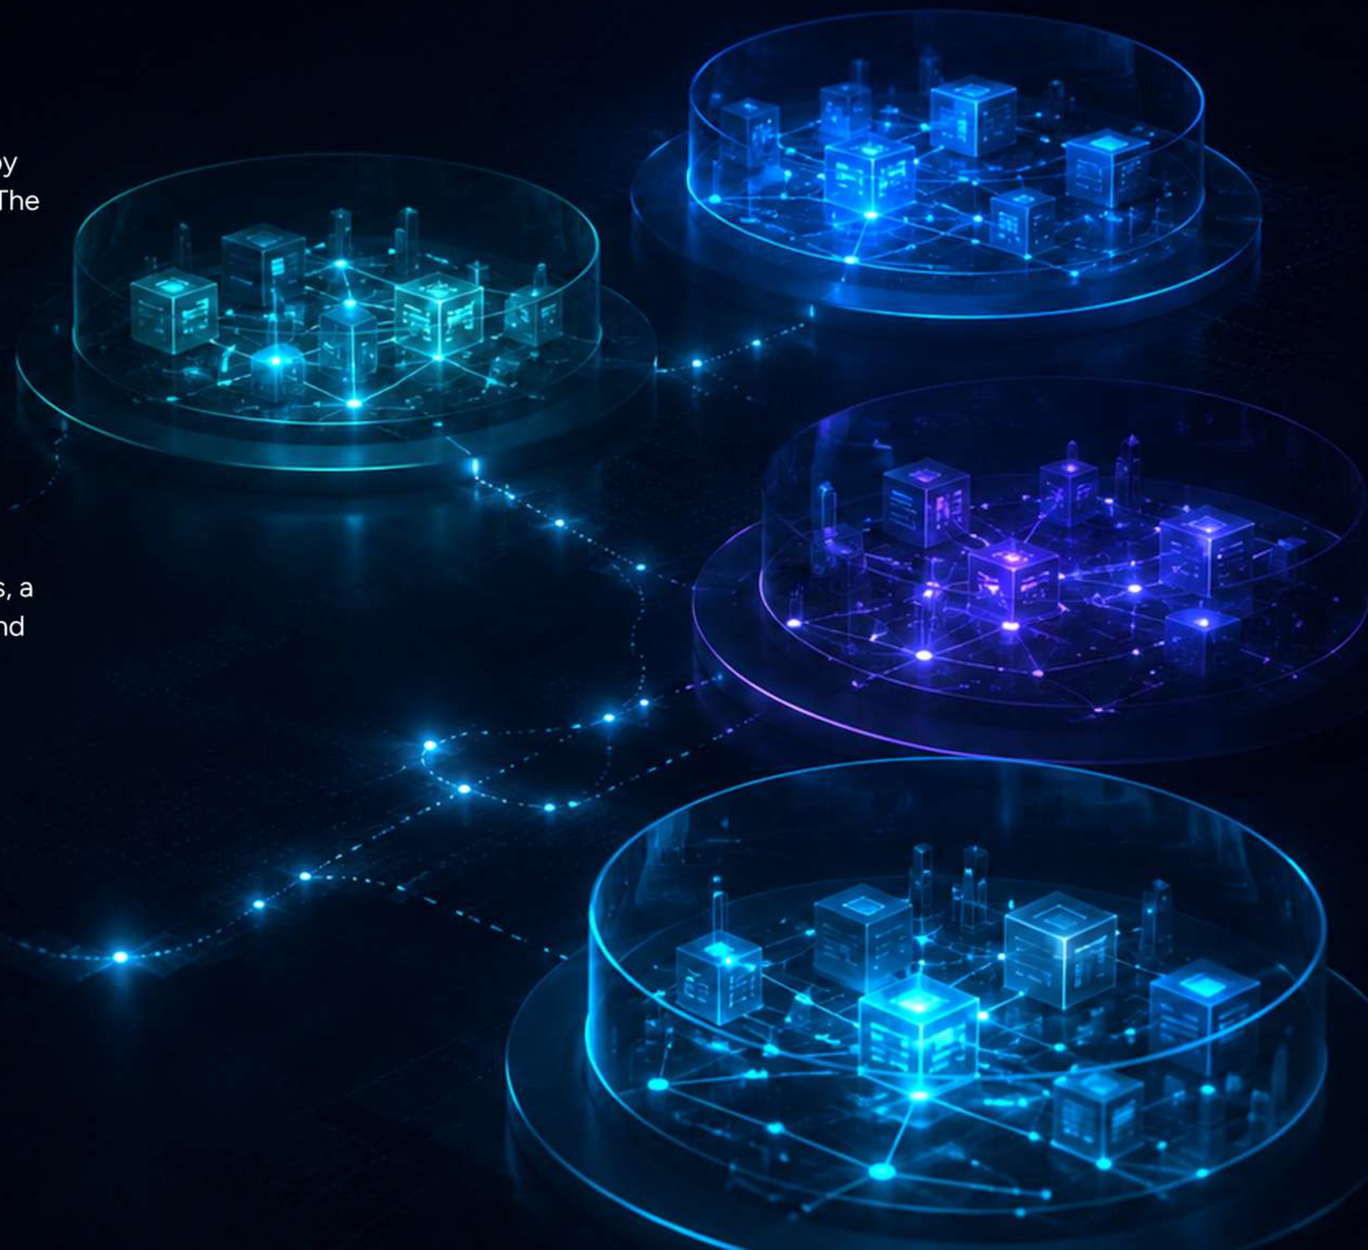

What the evidence shows

3 The accountability question is unanswered.

McKinsey finds that among organisations using AI, around half have already experienced at least one negative consequence. Yet, only about **1/3 have reached the higher maturity levels** for AI strategy, governance, and agentic-AI controls. And **78% of organisations have no formal policy** for creating or removing AI identities.

4 The control layer is maturing within platforms

Microsoft positions Agent 365 as the control plane for AI agents. AWS, Salesforce, and Google have converged on the same claim for Bedrock AgentCore, Agentforce, and Gemini Enterprise.

5 Regulators and risk leaders are pointing at the same gap.

Singapore's Model AI Governance Framework for Agentic AI states that humans remain ultimately accountable. The EU AI Act extends high-risk obligations across multi-agent chains from August 2026. Experian named **machine-to-machine activity its top fraud threat** for 2026, citing transactions with no clear owner of liability.

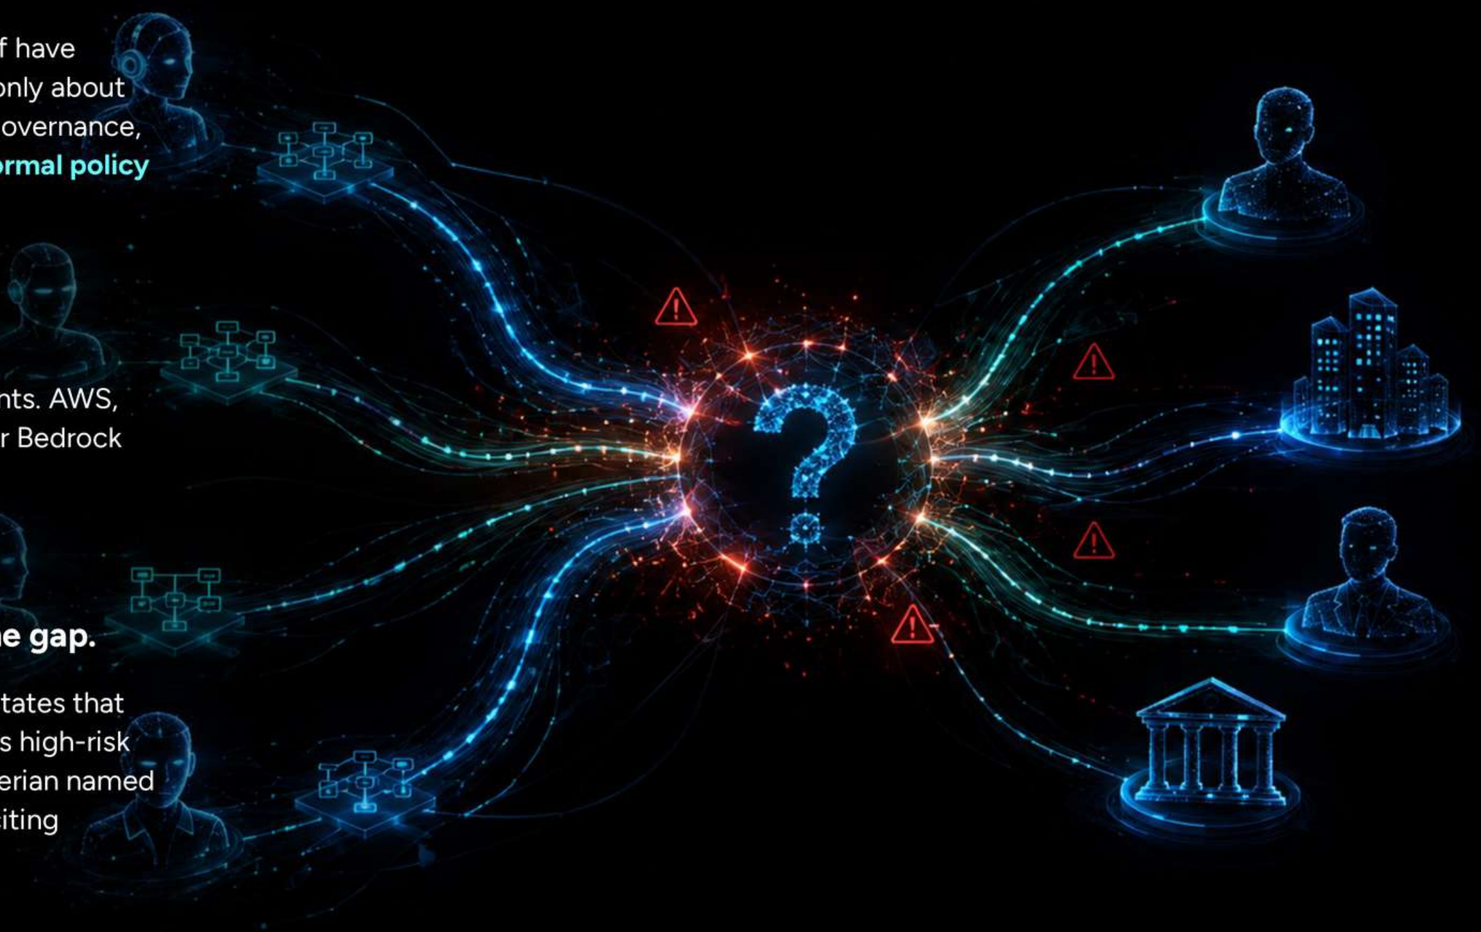

The market landscape: a strong consensus, and its blind spot

The agent-governance market is genuinely good at the thing it was built for, and any fair reading has to start there. The criticism in this report is narrow and specific, focusing on scope rather than quality.

What control planes do well

Control planes give an organisation a single place to deploy, observe, secure, and govern its AI agents within its own environment. The identity and risk specialists make the same core point Affinidi makes about human accountability.

SailPoint, for example, frames its agentic offering around mapping every AI agent back to a human owner. The principle is that each AI agent holds its own identity and stays bound to a human who answers for it, rather than borrowing that person's identity outright.

That expectation is now taking hold, a sign that the market is maturing in the right direction.

The blind spot they share

The same design choice that makes these control planes strong within a single platform is what prevents trust from leaving it. There is one property that every one of these systems has in common, and it is structural rather than a matter of effort or quality. Each of them roots an AI agent's identity in its own platform.

Microsoft's own documentation states that an Entra agent identity is single-tenant and cannot issue tokens outside the tenant that created it. Google's, Salesforce's, and AWS's identities are likewise platform-rooted.

That is exactly what an enterprise standardised on one platform wants. It is also the precise reason none of them could provide the same level of governance for an AI agent once it acts beyond its own environment.
This is not a flaw the incumbents can patch. Their identity model is rooted in their platform because their business is the platform. A portable identity that worked just as well off their cloud would weaken the thing that makes the control plane valuable to them. The boundary is built in.

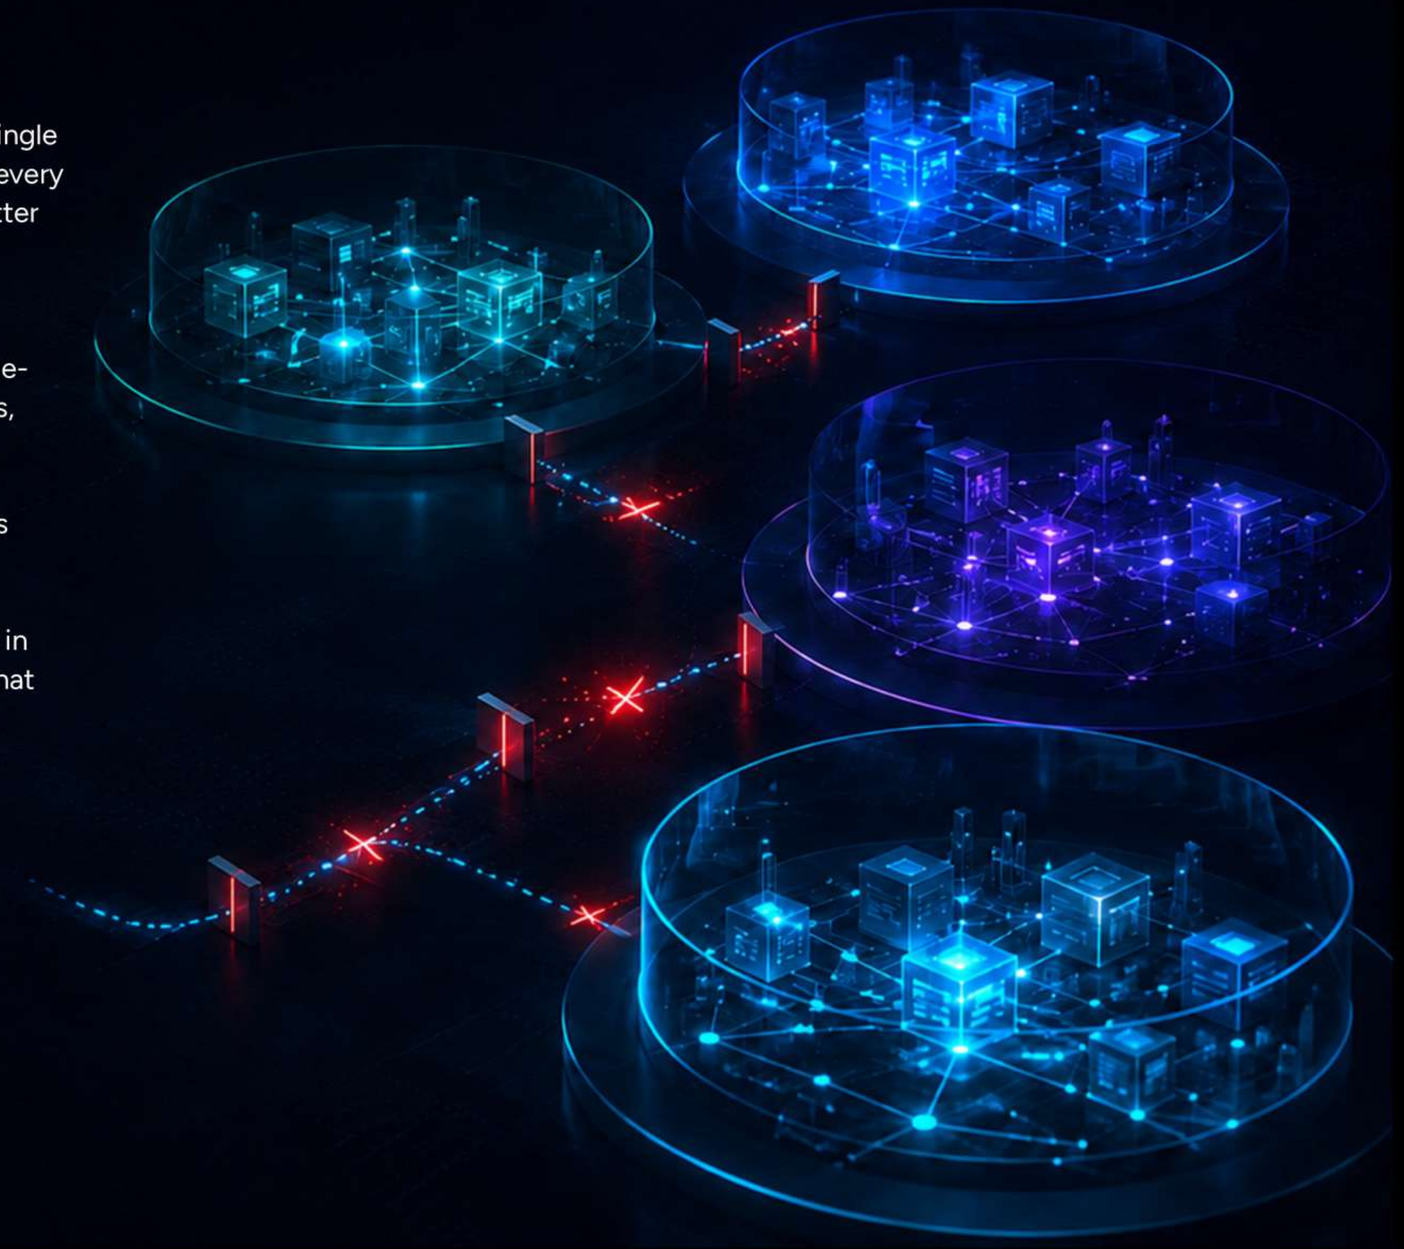

Continue from previous slide....

Salesforce roots native Agentforce agents in its organisation, and when it exposes external agents through the MuleSoft registry, it federates them to the enterprise's own directory, so the root still sits with a single provider that a counterparty must reach and trust.

The market has solved the governance problem within a single platform. Yet, it has barely begun to address what happens when an AI agent acts beyond that platform, whether the next step is another cloud, another service, or another organisation. For example, an AI agent built in Salesforce that calls a service like BambooHR crosses a boundary that Salesforce does not govern, even though both are connected within the company's architecture.

> **"This is where the whole agent-commerce industry is heading. Once you accept that trust cannot stop at your organisation's edge, the only question left is what carries it across, and that is what the rest of this report is about."**

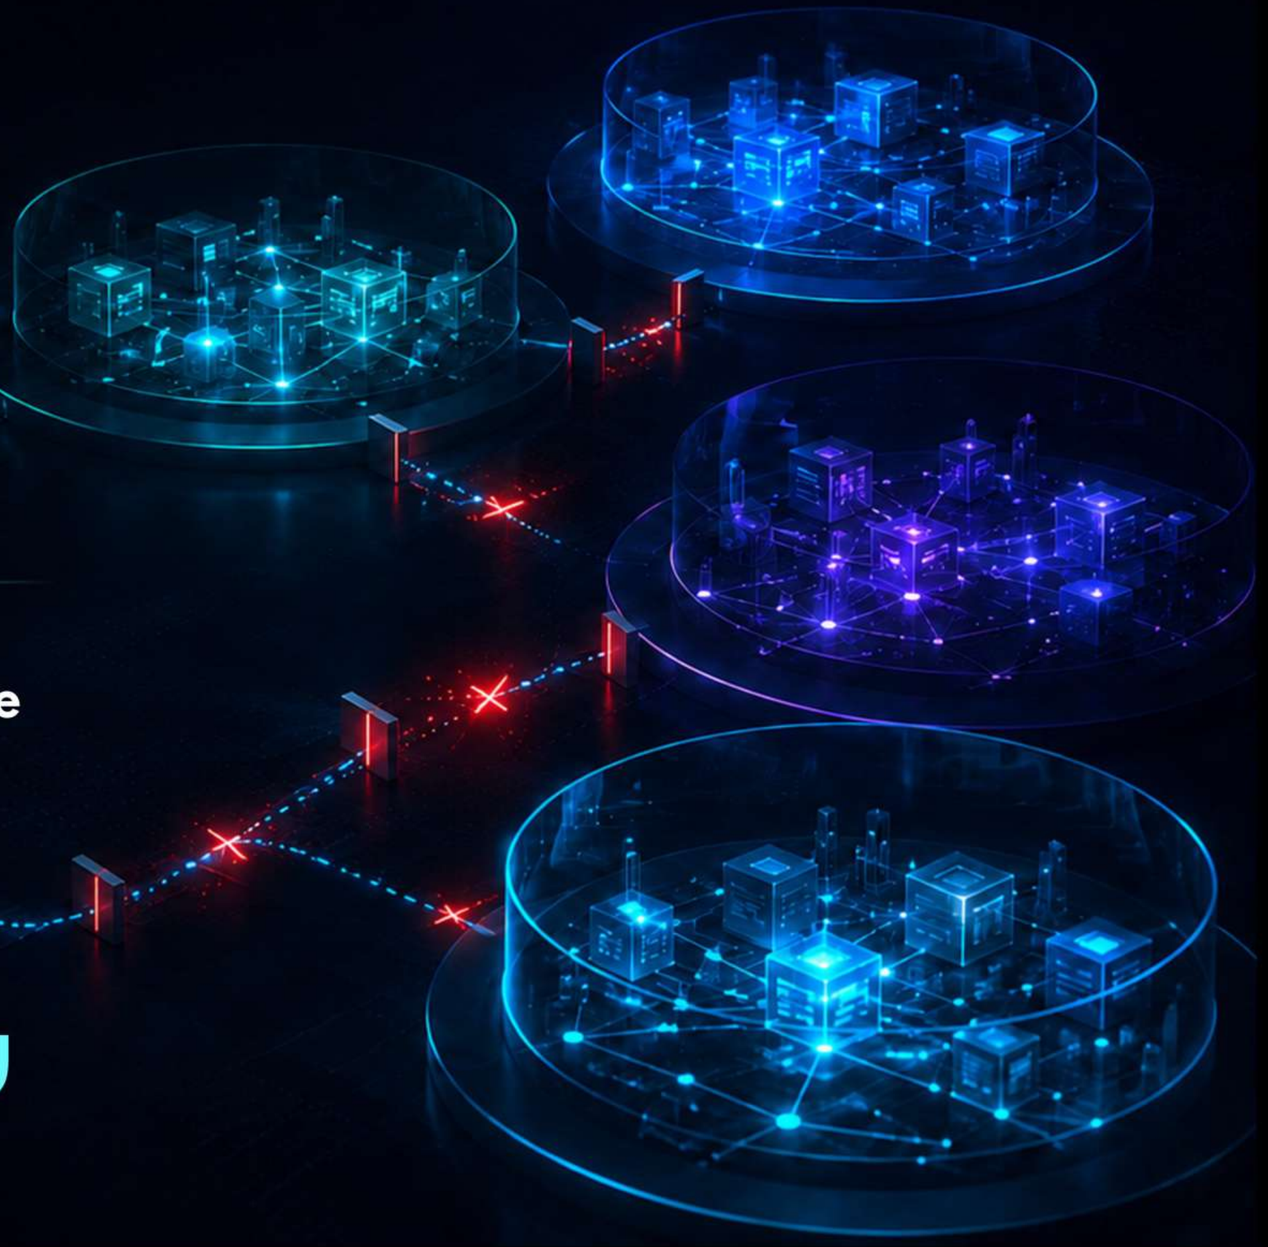

FINDINGS

Five findings, each drawn from independent evidence, build the case that the centre of gravity in AI agent governance is moving from control to portable accountability that **spans boundaries.**

Finding 1: Production is outpacing governance

AI Agents are being produced faster than the controls around them can catch up.

McKinsey's 2025 research found that among organisations using AI, 51% had already experienced at least one negative consequence, while just over a quarter (28%) reported that their CEO oversees AI governance. Fewer still reported that their board does. Its 2026 trust-maturity work found only about a third of organisations reaching a meaningful maturity level. Gartner expects more than 40% of agentic AI projects to be cancelled by the end of 2027, citing cost, unclear value and weak risk controls.

The constraint is trust maturity, so the projects that survive are those that can demonstrate portable accountability before they scale.

51%
of AI-using organisations have already hit a negative consequence (McKinsey, 2025)

~1 in 3
reach a meaningful level of AI governance maturity (McKinsey, Mar 2026)

40%+
of agentic AI projects cancelled by end 2027 (Gartner forecast)

Finding 2: The dominant control planes are single-tenant by design

The dominant control planes root AI agent identities in their own platforms, which serves inside the organisation and cannot enable provable trust beyond it.

As set out above, the dominant control planes root AI agent identity in their own platforms and govern within a single tenant. This is documented for Microsoft, whose Entra agent identity cannot issue tokens outside its tenant, and it holds by design for Google, which roots agent identity in a trust domain built on the open SPIFFE standard but anchored in Google's own managed platform, and for AWS, which issues workload identity rooted in AWS. Salesforce roots native Agentforce agents in the Salesforce org. Where it opens to external agents via the MuleSoft registry, it federates to the enterprise's OIDC provider and enforces trust at the MuleSoft gateway rather than issuing identities itself.

There are several gaps to solve for cross-boundary agentic workflows. The agent identity is still bound to the platform and not portable, or cannot be verified independently without accessing the root platform. Existing control planes are structurally unable to prove verifiable trust outside their perimeters yet. Unlike human interaction, agentic collaboration, and even autonomous agents, which need verifiable proof to be trusted, this conflicts with many tech hyperscalers' business models of closed, walled gardens.

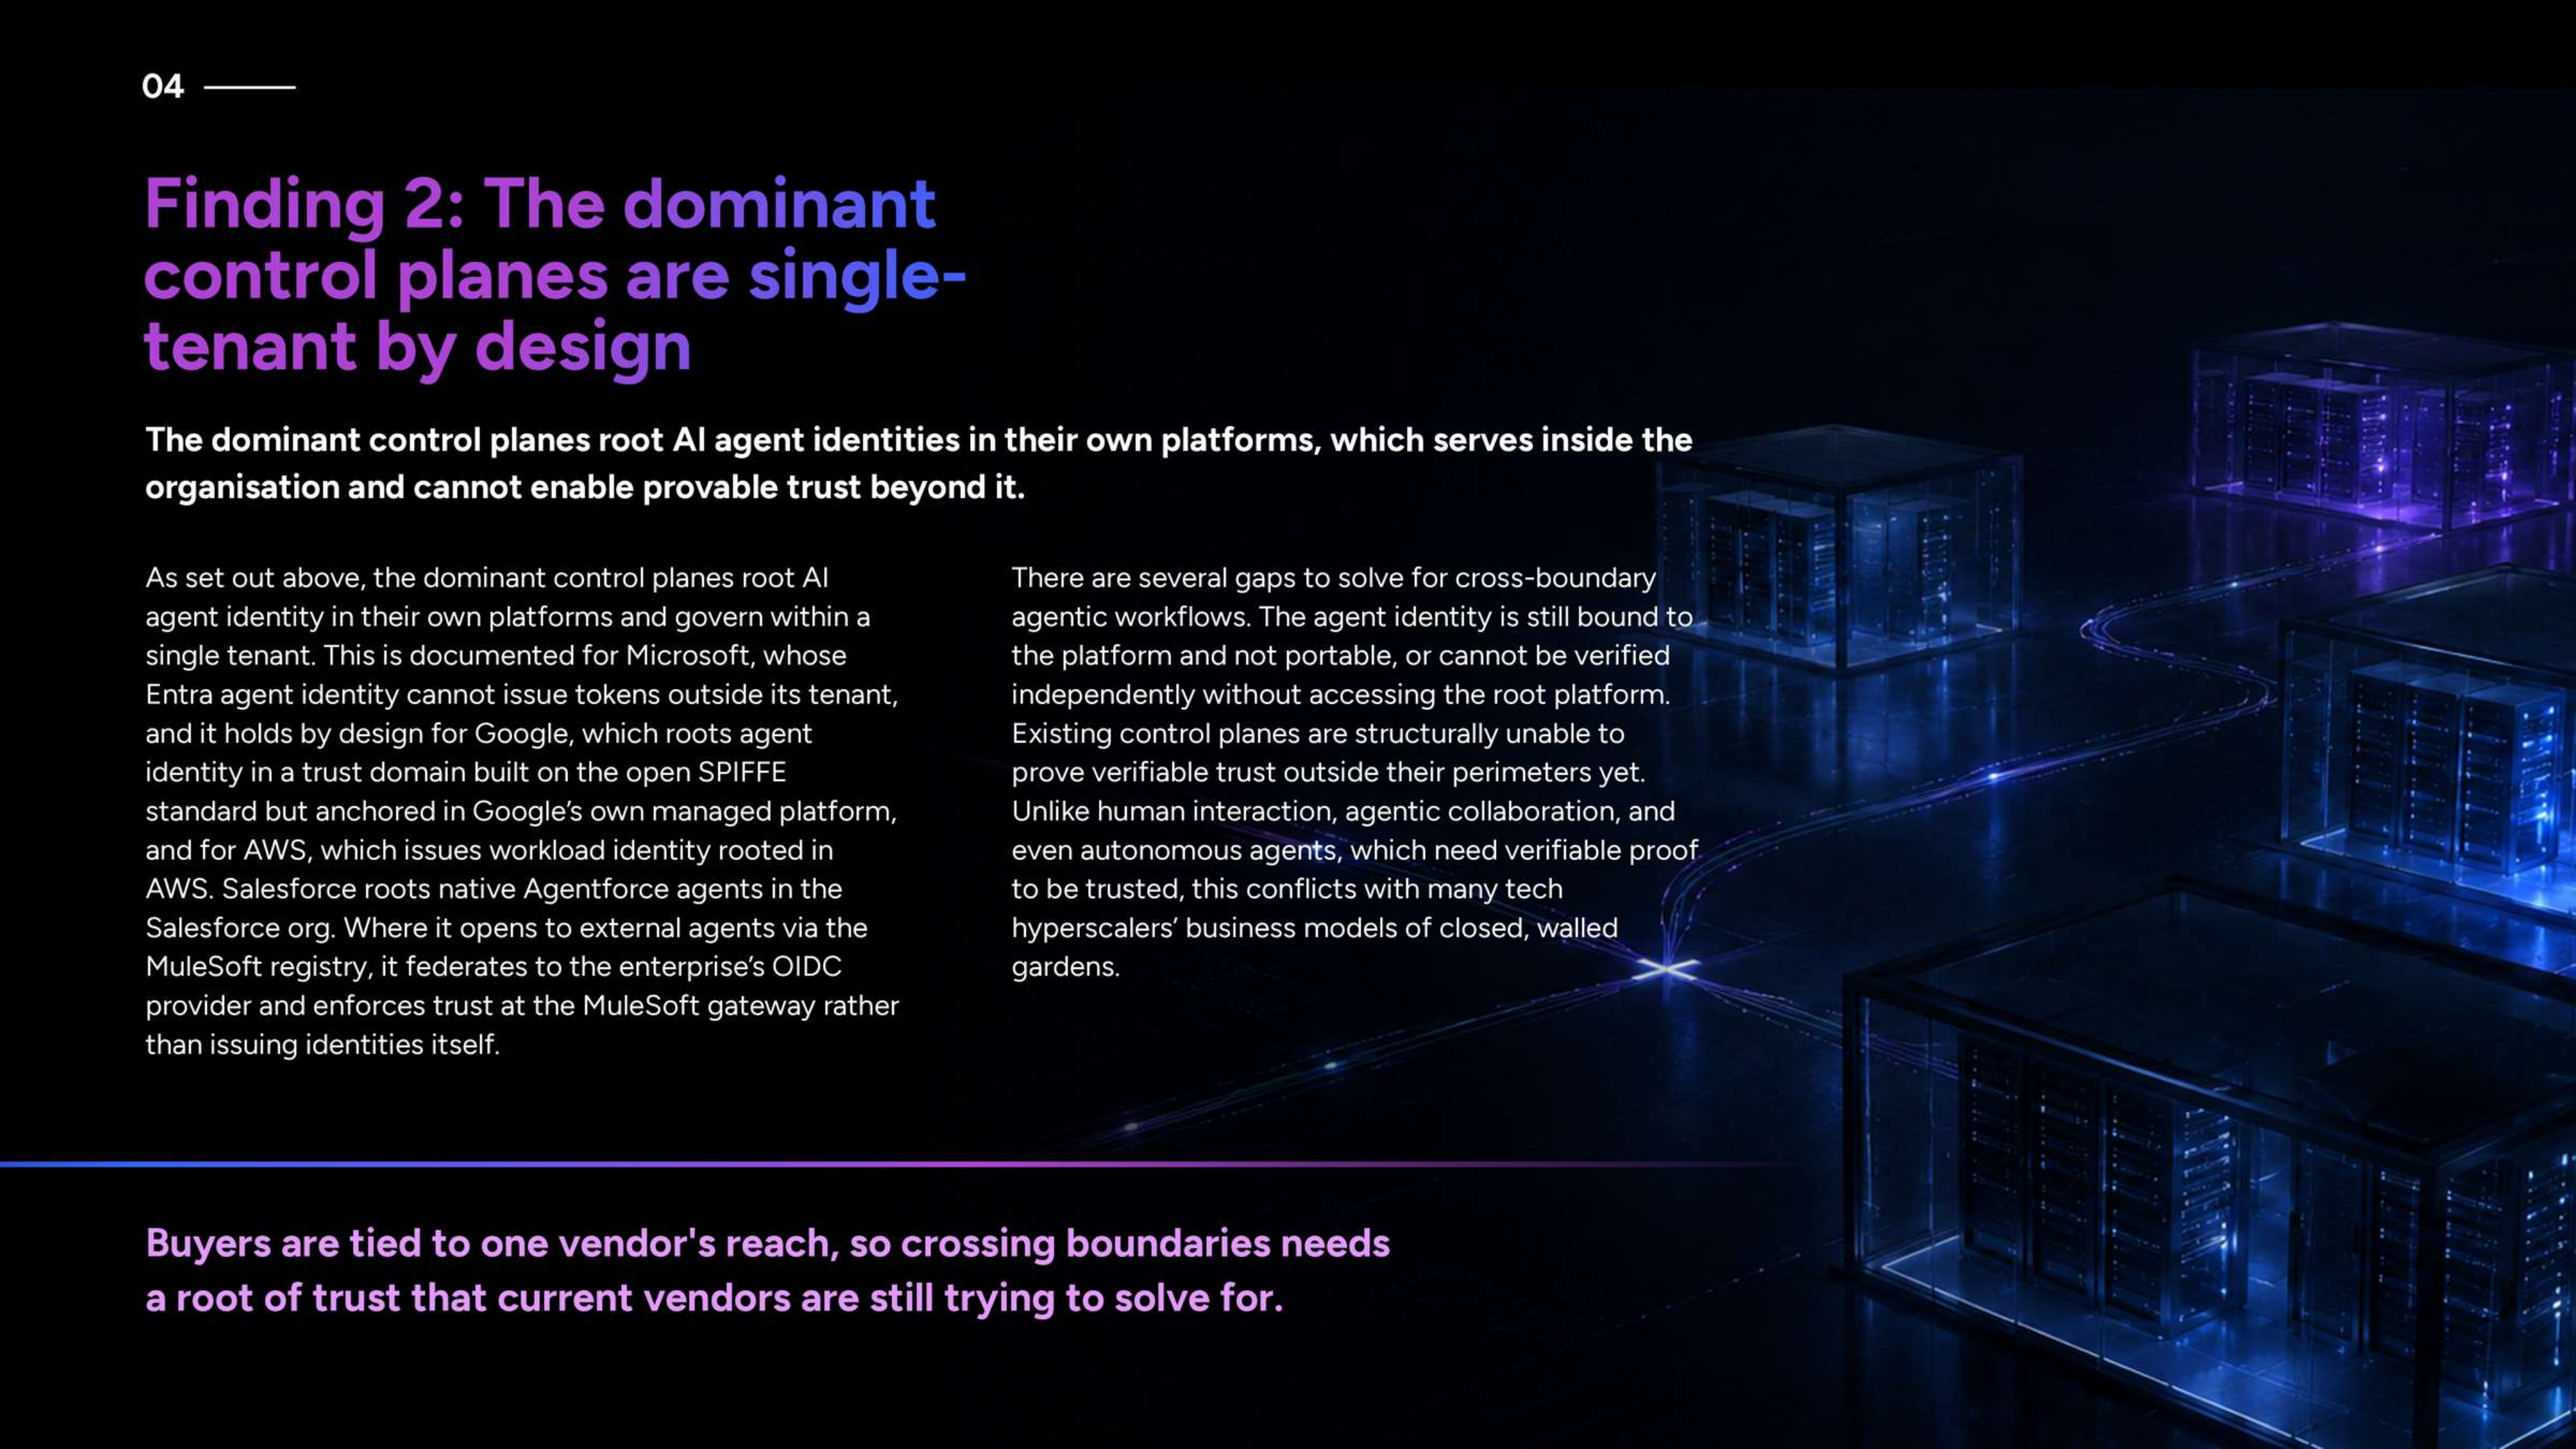

Buyers are tied to one vendor's reach, so crossing boundaries needs a root of trust that current vendors are still trying to solve for.

Finding 3: Accountability is the layer nobody owns

Regulators are converging on one question about AI agents: once an AI agent acts beyond one's walls, who answers for it?

Singapore has drawn the clearest line. In January 2026 it published what it calls a world-first governance framework for agentic AI, and its central rule is blunt: a human stays accountable for what an agent does, and every agent is tied to a named supervisor. That is the hero fact of this finding, because it assigns accountability to a person rather than a platform, and does so through law-adjacent guidance rather than vendor's marketing.

Two further signals point the same way.

From August 2026, the EU AI Act extends its high-risk duties across every agent in a multi-agent chain, so accountability no longer stops at one's own agent. And Experian's 2026 fraud forecast names machine-to-machine activity its top threat, warning of transactions fired off with no clear owner of liability. Singapore says who must answer, the EU says the duty travels down the chain, Experian shows the cost when nobody does.

Jan 2026

Singapore publishes a world-first agentic-AI governance framework tying every agent to a supervising human (IMDA).

Aug 2026

EU AI Act high-risk duties extend across every AI agent in a multi-agent chain (Reg. 2024/1689).

2026's Top Threat

Experian names machine-to-machine activity the top fraud threat of the year.

The regulatory line is moving in one direction: accountability should stay attached to a human and the agent as it crosses boundaries.

Finding 3: Accountability is the layer nobody owns

China has moved the same way, and it has released the standard for "Artificial Intelligence Agent Interconnection", the country's first national standard on agent connectivity, issued by its State Administration for Market Regulation, setting out a unified identity management framework for secure interaction between AI agents across platforms and systems.

Regulators and risk leaders have stopped asking whether an AI agent can be watched. They are now asking who answers for it once it crosses a boundary. Accountability needs proof, not just records and logs. Watching an AI agent inside your own systems tells you nothing about who is accountable once it acts somewhere else.

> **When an agent of ours settles with an AI agent we have never met, the question our regulator will ask is simple: who authorised this, and who is liable. We need that answer to travel with the transaction, not sit in someone else's platform.**

Financial services, Hyman Zu, Cofounder of CardInfoLink

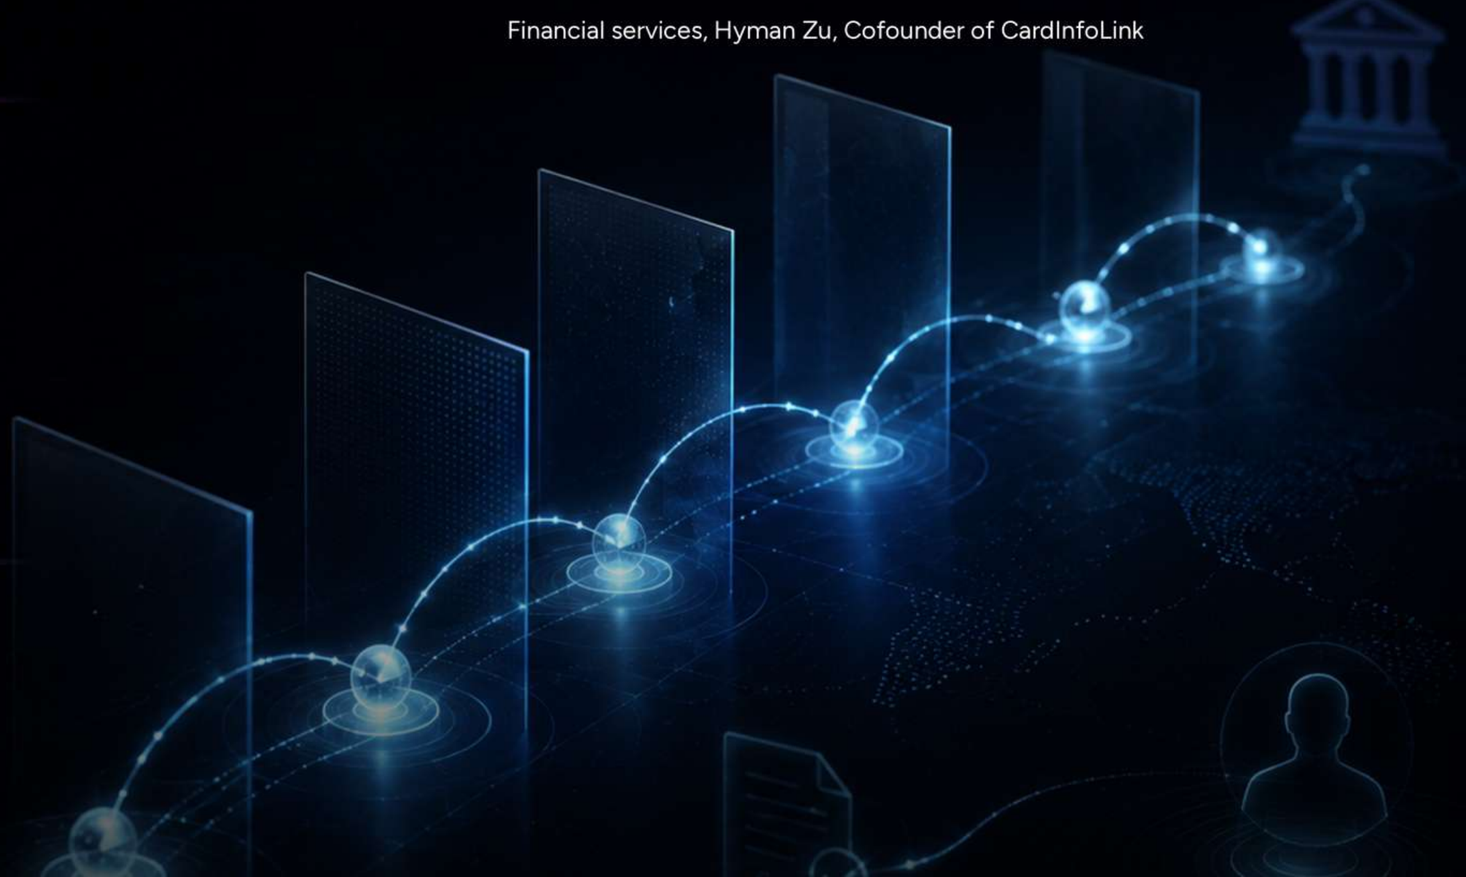

Finding 4: The identity surface is exploding and ungoverned

The number of identities to account for is growing far faster than the policies to manage them.

Vendor research consistently reports machine identities vastly outnumbering human ones, with Palo Alto Networks putting the ratio at roughly **109 to 1 in 2026**, up from 82 to 1 a year earlier, and around **79 of every 109** being AI agents.

These figures are self-reported and vary widely by source and environment, with other vendors citing ratios from 45 to 1 up to 144 to 1, so they are best read as directional. The governance response lags badly: the Cloud Security Alliance found that fewer than a quarter of organisations have a formal policy for creating or removing AI identities.

The sheer volume renders directory-based control ineffective, so identity must be agile and verifiable by parties outside the issuer.

~109:1

machine-to-human identity ratio in 2026, up from 82:1 a year earlier (Palo Alto; range 45:1 to 144:1 across vendors)

‹25%

of organisations have a formal policy for creating or removing AI identities (Cloud Security Alliance)

Finding 5: The market is already building portable trust out of necessity

The cross-boundary rails are the market voting with engineering effort for portable, vendor-neutral trust.

A2A, AP2, and Mastercard's and Visa's agent-commerce products each had to create a way for an AI agent to present a verifiable identity and authority to a counterparty on another platform. Google's decision to donate AP2 to the FIDO Alliance and the Linux Foundation's neutral stewardship of A2A show that even the largest players accept that these foundations cannot be owned by any single vendor.

The rails already exist, so the open prize is the identity layer that rides them.

150+

supporting organisations behind A2A under neutral Linux Foundation governance

100+

partners on AP2, now donated to the FIDO Alliance

WHERE THE MARKET IS HEADING

The enterprise story: how do you make AI agents your moat

Companies are rapidly jumping into the AI agent transformation journey. The ability to build, train, and deploy AI agents is already a commodity. The ability to unleash and scale them without increasing liability is the new moat. The findings point in one direction: agent governance is moving from control rooted in a single platform towards trust portable across platforms.

The Enterprise AI Maturity Curve Diagram

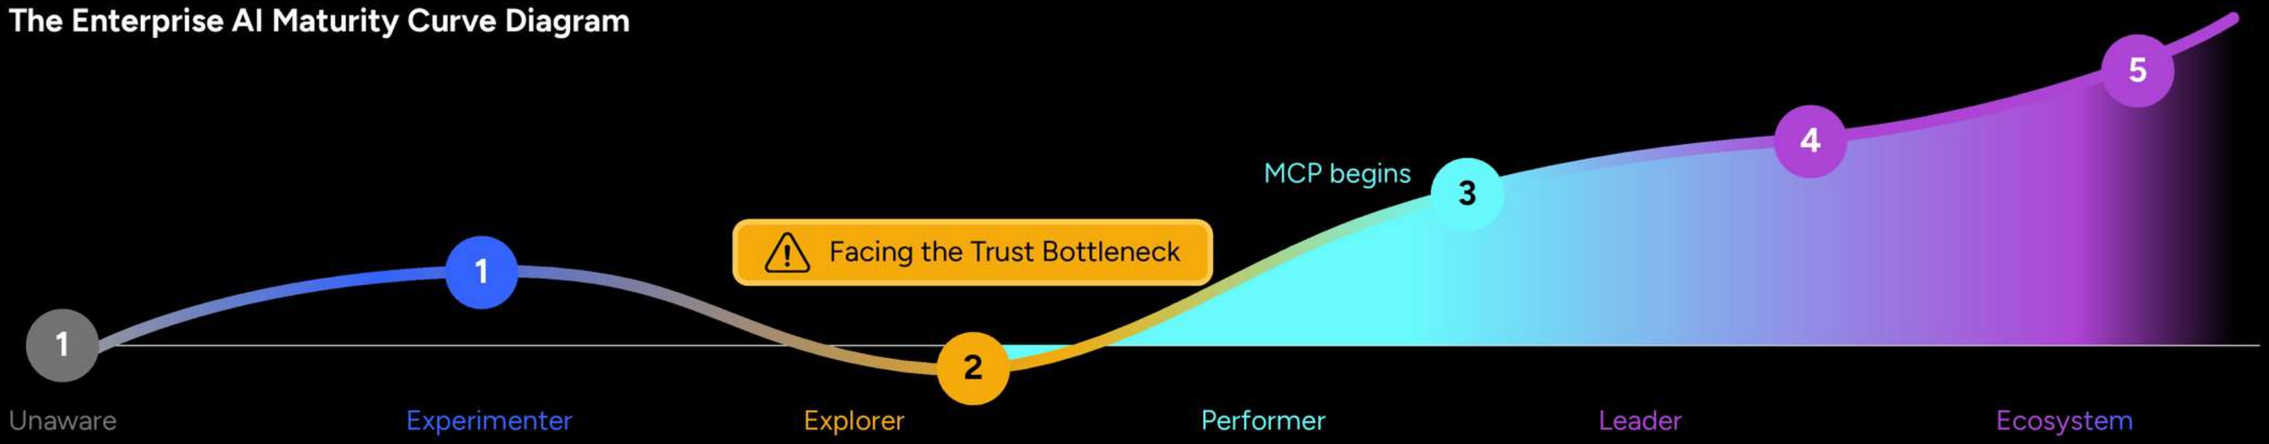

1

- AI pilots and POCs
- Use-case specific

2

- Isolated AIs deployed in production
- Human-in-the-loop

3

- AI workflow cross trust-boundary
- Agent-to-Agent communication without human-in-the-loop

4

- Multi-agent systems in production
- Agent-to-agent sharing sensitive information

5

- Cross-org use-cases
- Agents interacting with unknown agents

Figure 1. The enterprise AI maturity curve is a typical path of agentic AI adoption from internal pilots to cross-organisation ecosystems. The trust bottleneck marks the point at which agents begin to act beyond a single platform's control, and it is here that portable identity and accountability become the constraints on further progress. Curve shape and stage markers are illustrative.

WHERE THE MARKET IS HEADING

The enterprise story: how do you make AI agents your moat

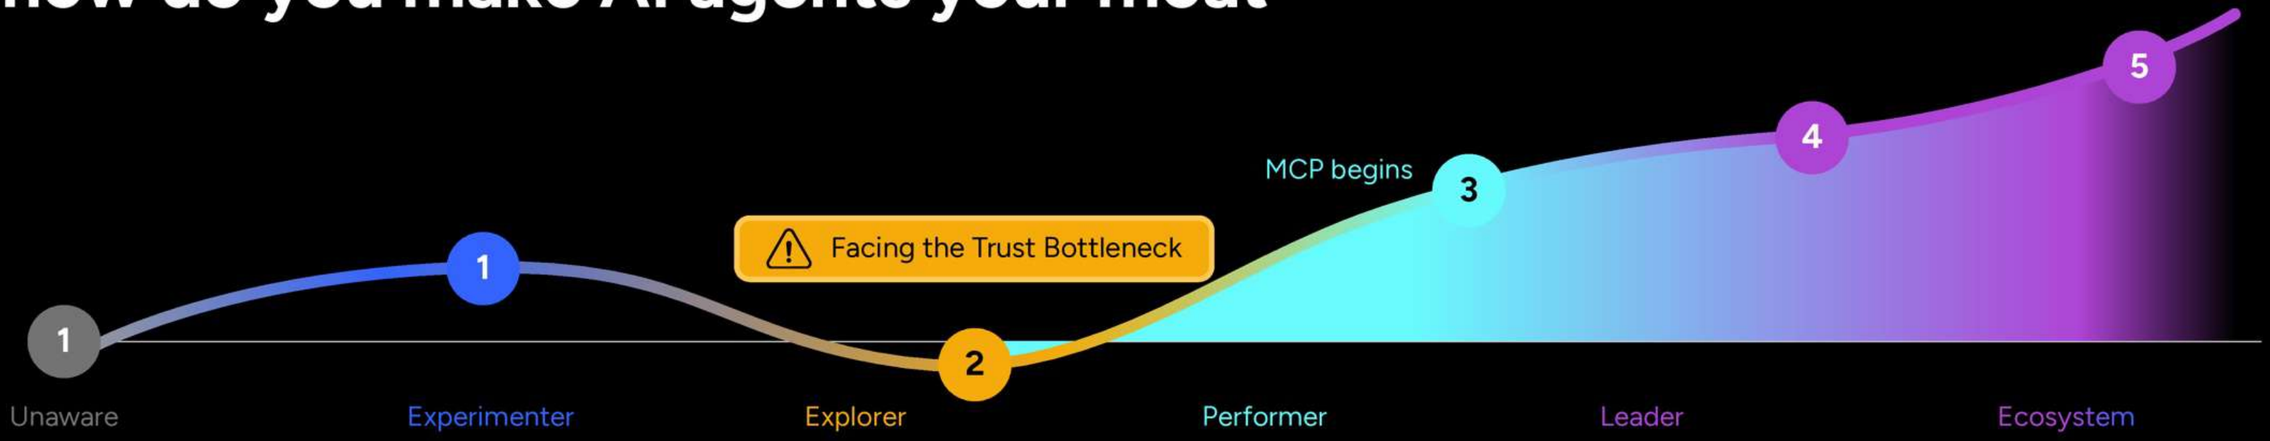

Reading the curve. The path runs from internal pilots to cross-organisation ecosystems, and the trust bottleneck sits between the two.

Stage 0

Unaware:
no agent activity yet.

Stage 1

Experimenter:
pilots and proofs of concept, use-case specific.

Stage 2

Explorer:
isolated agents in production with a human in the loop, where the trust bottleneck begins.

Stage 3

Performer:
agent workflows that cross trust boundaries, and agent-to-agent interaction without a human in the loop, where cross-boundary rails such as MCP, A2A, and AP2 come into play.

Stage 4

Leader:
multi-agent systems in production, with agents exchanging sensitive information.

Stage 5

Ecosystem:
cross-organisation use cases, with agents transacting with agents they have never met.

The competitive frontier is the extent to which an enterprise can push AI agent value in its business. The strategic advantage will be the ability to scale AI agents across platforms, clouds, and organisations, while maintaining trust and avoiding liability.

Protocol adoption: how fast the market is moving

If the destination is where the market is heading, the speed of travel is measured by the rate of protocol adoption. These protocols are the rails the market is laying, and Affinidi's identity layer is built to ride whichever of them takes hold, so faster protocol adoption widens the ground a portable identity layer has to serve. The table below is a desk-research baseline of the major cross-boundary protocols, the kind of proprietary cut that, maintained over time, becomes a citable signal in its own right.

PROTOCOL	GOVERNANCE	PURPOSE	ADOPTION SIGNAL (DATED)	SOURCE TO ADD
A2A (Agent2Agent)	Linux Foundation (neutral)	Agent-to-agent interoperability across platforms	50+ to 150+ supporting organisations between Apr 2025 and Apr 2026; v1.0 adds signed agent cards	Google Developers Blog (9 Apr 2025, launch/50+); Linux Foundation press release (23 Jun 2025, contribution); Linux Foundation press release "A2A Protocol Surpasses 150 Organisations" (9 Apr 2026, 150+); A2A Protocol Specification / "Announcing Version 1.0" at a2a-protocol.org (signed Agent Cards).
AP2 (Agent Payments Protocol)	Donated to FIDO Alliance	Agent-to-merchant payments	60+ partners at launch (Sep 2025), passed 100 by late Oct 2025; v0.2 adds human-not-present payments	Google Cloud blog "Announcing Agent Payments Protocol (AP2)" (16 Sep 2025, 60+); Google Cloud blog "Introducing an agentic commerce solution for merchants from PayPal and Google Cloud" (c. 27 Oct 2025, 100+); Google blog.google (28 Apr 2026, FIDO donation + v0.2 / Human Not Present); AP2 specification at ap2-protocol.org and google-agentic-commerce/AP2 GitHub (technical detail).
MCP (Model Context Protocol)	Linux Foundation (Agentic AI Foundation)	Agent-to-tool context and actions	More than 10,000 active public MCP servers and 97M+ monthly SDK downloads (Python + TypeScript), per Anthropic and the Linux Foundation, 9 Dec 2025. Created by Anthropic (Nov 2024); donated to the Linux Foundation's Agentic AI Foundation 9 Dec 2025. Next specification revision (2026-07-28) moves MCP to a stateless core with an extensions framework; release candidate locked 21 May 2026, final publication scheduled 28 July 2026.	"Donating the Model Context Protocol and establishing the Agentic AI Foundation" (9 Dec 2025); Linux Foundation press release "Linux Foundation Announces the Formation of the Agentic AI Foundation (AAIF)" (9 Dec 2025); official MCP blog "MCP joins the Agentic AI Foundation" (9 Dec 2025). Official MCP blog "The 2026-07-28 MCP Specification Release Candidate" (21 May 2026).

PROTOCOL	GOVERNANCE	PURPOSE	ADOPTION SIGNAL (DATED)	SOURCE TO ADD
Mastercard Agent Pay for Machines	Mastercard	Verifiable agent identity and multi-rail settlement	30+ participants at launch (Jun 2026)	Mastercard press release "Mastercard Unveils Agent Pay" (29 Apr 2025, launch); PYMNTS "PayOS, Mastercard Complete First AI Payment" (29 Sep 2025, first live agentic transaction); Mastercard Newsroom, live agentic transactions in Hong Kong (27 Mar 2026) and Thailand (7 Apr 2026); Mastercard press release "Mastercard Launches Agent Pay for Machines" (10 Jun 2026, 30+).
Visa Intelligent Commerce	Visa	Agent-initiated commerce across multiple agent protocols	More than 100 partners globally, with 30+ building in the VIC sandbox and 20+ agents/enablers integrating directly (Visa, 18 Dec 2025)	Visa press release "Find and Buy with AI: Visa Unveils New Era of Commerce" (30 Apr 2025, launch); Visa press release "Visa and Partners Complete Secure AI Transactions" (18 Dec 2025); Visa press release "Visa Opens the Door to AI-Driven Shopping for Businesses Worldwide" (8 Apr 2026, Intelligent Commerce Connect pilot partners). Visa Payments Forum 2026 announcement (Jun 2026).
UCP (Universal Commerce Protocol)	Vendor-steered Tech Council (Google and Shopify with commerce incumbents); open standard	Agent-led commerce orchestration across discovery, cart, checkout, and post-purchase	Launched January 2026 with 20+ endorsers; Tech Council expanded to ten members on 24 Apr 2026 as Amazon, Meta, Microsoft, Salesforce, and Stripe joined founders Google, Shopify, Etsy, Target, and Wayfair; Shopify made UCP self-serve from 17 Jun 2026	Google UCP launch announcement (Jan 2026); UCP Tech Council expansion announcement (24 Apr 2026); Shopify Summer '26 Edition / Developer Dashboard (Jun 2026).
ACP (Agentic Commerce Protocol)	OpenAI and Stripe; open source (Apache 2.0)	In-surface checkout execution between agents and merchants, with merchant of record retained	Launched 29 Sep 2025 alongside ChatGPT Instant Checkout; broad US "Buy it in ChatGPT" 16 Feb 2026; consumer Instant Checkout product retired March 2026 while the protocol continues; over 1M Shopify merchants enrolled	Stripe and OpenAI ACP announcement (29 Sep 2025); ACP specification at agenticcommerce.dev; OpenAI ChatGPT shopping update (Mar 2026).

PROTOCOL	GOVERNANCE	PURPOSE	ADOPTION SIGNAL (DATED)	SOURCE TO ADD
x402	Linux Foundation (x402 Foundation); neutral	HTTP-native machine-to-machine payment settlement in stablecoins, reviving the dormant HTTP 402 status code	Created by Coinbase (May 2025); contributed to the Linux Foundation x402 Foundation on 2 Apr 2026 with 20+ founding members including AWS, Circle, Cloudflare, Google, Mastercard, Microsoft, Stripe, and Visa; Coinbase reports 169M+ payments in the first year; AWS CloudFront and Cloudflare edge integrations Jul 2026	Linux Foundation press release "Launching the x402 Foundation" (2 Apr 2026); Coinbase x402 announcement (May 2025); Coinbase first-year metrics (Jul 2026).
Web Bot Auth	IETF (Web Bot Auth Working Group); Cloudflare-led	Cryptographic verification of automated agent traffic via signed HTTP requests (RFC 9421)	Cloudflare proposal (2025); open key-registry format with Amazon Bedrock AgentCore Feb 2026; IETF Working Group chartered 2026; verifying agents include Anthropic, OpenAI, and Perplexity; implemented by AWS, Akamai, Vercel, and Shopify	IETF draft-meunier-web-bot-auth-architecture; Cloudflare Web Bot Auth / verified-bots documentation; IETF Web Bot Auth Working Group charter (2026).

"Adoption of cross-boundary protocols is compounding quarter on quarter and consolidating under neutral governance. The rails for agent-to-agent commerce are being laid in the open. What is not yet settled is the identity layer that lets an AI agent prove itself and its human principal as it rides them."

OUTLOOK: FIVE PREDICTIONS FOR 2026 TO 2028

01 Cross-boundary agent liability becomes a board-level issue by 2027, starting with regulated industries

With the EU AI Act in effect from August 2026 and Experian naming machine-to-machine activity a top fraud threat for 2026, accountability for what an AI agent does across organisations shifts from an IT topic to a board-level risk. Regulated sectors such as financial services and healthcare will prioritise it first, since they combine the tightest compliance exposure with the earliest cross-boundary agent use

2027

What would change our mind:
a coordinated industry liability framework lands earlier and defuses the question.

OUTLOOK: FIVE PREDICTIONS FOR 2026 TO 2028

02 By 2028, most enterprises running agents in production will require a portable identity for any AI agent acting outside their own environment.

The cross-boundary rails already assume it. Protocols such as A2A and AP2 require an agent to prove its identity and the identity of the party standing behind it to a counterparty outside its own boundaries (platform, cloud, or system).

2028

What would change our mind:
one vendor's identity becomes a de facto cross-tenant standard that the market simply accepts.

03 At least one major control-plane vendor adopts or acquires portable, cross-tenant agent credentials before the end of 2027.

These credentials allow an AI agent to prove its identity across platforms. Okta has already signalled support for verifiable credentials for FY27, and the wave of identity-security acquisitions points the same way. The capability already exists in the market, and the pattern to watch is the incumbents being pulled towards it.

2027

What would change our mind:
incumbents judge that defending platform lock-in is worth more than closing the gap.

OUTLOOK: FIVE PREDICTIONS FOR 2026 TO 2028

04 “Who is the human principal?” becomes a standard procurement question for AI agent platforms by 2027.

Driven by multi-agent accountability rules in the EU AI Act and frameworks such as Singapore’s.

2027

What would change our mind:
regulators accept organisation-level rather than human-principal accountability.

05 The agentic projects that survive the coming cancellations are disproportionately interoperable.

Gartner expects more than 40% of agentic AI projects to be cancelled by the end of 2027, and single-vendor deployments that cannot cross boundaries are the most exposed to that cost-and-value squeeze.

What would change our mind:
cancellations cluster instead around immature use cases, regardless of architecture.

What this means for the trust layer

The agentic economy needs something that current offerings were never designed to provide: an identity that an AI agent carries, proving who it is and who is accountable for it to any party on any platform, and that keeps working when the AI agent changes cloud, model, or vendor. That is a different layer from the control plane. It sits beneath cross-organisation transactions rather than within one estate.

This is the layer Affinidi builds.

Agent Gateway, a part of the Affinidi Trust Fabric suite, gives each AI agent a portable, verifiable identity rooted in open standards (W3C Decentralised Identifiers and Verifiable Credentials) rather than in any single vendor's tenant, with cross-boundary trust tunnelling and a human principal bound in.

It is designed for exactly the destination the framework describes: vendor-neutral and built to cross boundaries, complementing control planes and enabling portable trust across clouds, platforms and organisations.

To find out more about **Agent Gateway: click here**

> **"An AI agent is only as trusted as the identity it can prove."**

Methodology and evidence grading

This is a point-of-view report on the category and the market. The analytical core rests on independent, third-party evidence, and the whole field is named and characterised fairly, including the players whose architecture this report treats as a foil, as well as Affinidi itself.

1 Evidence base

Every external claim is graded, in Appendix A, as one of three types. Verified means a primary regulatory source, a dated corporate announcement or independent reporting. Forecast means an analyst or vendor projection, phrased as a prediction and never as a measured outcome. Vendor survey means a self-reported figure from a commercially interested party, attributed to its source and treated as directional.

2 Currency

Fast-moving counts (protocol partner numbers, installed-base figures, machine-identity ratios) are dated to the month and were current as of June 2026. Partner counts in particular move quickly and should be re-checked at publication.

3 Numbers as ranges

Market-size estimates for agentic AI vary by an order of magnitude across research houses. Where a figure is contested, this report gives a range and attributes it rather than leaning on a single headline number.

4 Proprietary signal

Two analyses are presented as desk-research baselines: the enterprise AI maturity curve diagram and the protocol-adoption tracker.

5 Vendor architecture

Vendor-architecture claims are drawn from each vendor’s own documentation where it exists, and desk-assessed where it does not. The Salesforce characterisation is documented for native Agentforce agents and desk-assessed for external agents in the MuleSoft registry

Appendix A: Evidence based, graded

Key external claims used in this report, with source, date and grade. Grades: Verified (primary or independently corroborated), Forecast (analyst or vendor projection), Vendor survey (self-reported, directional).

CLAIM	SOURCE	DATE	GRADE
Over 40% of agentic AI projects cancelled by end 2027	Gartner	Jun 2025	Forecast
Guardian agents 10–15% of agentic AI market by 2030	Gartner	Jun 2025	Forecast
51% of AI-using organisations saw at least one negative consequence	McKinsey, State of AI in 2025 (Nov 2025)	2025	Verified (survey)
~A third reach meaningful AI governance maturity	McKinsey trust survey	Mar 2026	Verified (survey)
Machine-to-machine named top fraud threat for 2026	Experian Fraud Forecast	Jan 2026	Forecast / vendor
US consumers lost over $12.5bn to fraud in 2024	FTC (via Experian) 2024	2024	Verified
Humans remain ultimately accountable for agents	Singapore MGF for Agentic AI (IMDA)	Jan 2026	Verified
High-risk obligations apply across multi-agent chains	EU AI Act (Reg. 2024/1689)	From Aug 2026	Verified
A2A passes 150+ supporting organisations	Linux Foundation	Apr 2026	Verified
AP2 passes 100 partners (by Oct 2025); donated to FIDO Alliance (Apr 2026)	Google	2025–26	Verified (announcement)
Mastercard Agent Pay for Machines, 30+ participants	Mastercard	Jun 2026	Verified (announcement)
Visa Intelligent Commerce, 100+ partners; consumer product in pilot as of Q1 2026	Visa	2025–26	Verified (announcement)
UCP Tech Council expands to ten members; agent-led commerce orchestration launched Jan 2026	Google / UCP Tech Council	2026	Verified (announcement)
ACP launched Sep 2025; consumer Instant Checkout retired Mar 2026, protocol continues	OpenAI and Stripe	2025-26	Verified (announcement)
x402 contributed to the Linux Foundation x402 Foundation; 169M+ payments in first year	Coinbase / Linux Foundation	2026	Verified (announcement)
Web Bot Auth verifies agent traffic via signed HTTP requests; IETF Working Group chartered	IETF / Cloudflare	2026	Verified (standard)
Machine-to-human identity ratio ~109:1 (45:1 to 144:1)	Palo Alto Networks (109:1, 2026); CyberArk (82:1, 2025); Gartner (~45:1); Entro Security (144:1, contested with Oasis Security)	2025-26	Vendor survey
<25% of organisations have formal AI-identity policy	Cloud Security Alliance	Jan 2026	Verified (survey)
Agent 365 is "the control plane for AI agents"; GA at $15/user/mo	Microsoft	May 2026	Verified (vendor)
Agentic AI governance market $7.28bn (2025) to $38.94bn (2030)	Mordor Intelligence	2025	Forecast (single house)
Entra agent identity is single-tenant, cannot issue tokens outside its tenant	Microsoft Learn	2026	Verified (vendor doc)
Agent identity rooted in a Google-managed SPIFFE trust domain	Google Cloud IAM docs	2026	Verified (vendor doc)
Workload identity rooted in AWS; tokens for first-party services only	AWS Bedrock AgentCore docs	2026	Verified (vendor doc)
Native Agentforce agents rooted in the Salesforce org; external registry agents anchored in the enterprise's own OIDC provider via MuleSoft	Salesforce Engineering blog; MuleSoft docs	2026	Verified (native) / desk-assessed (external)

affinidi

Thank you

To learn more about Agent Gateway, click here

AUTHOR

Brian Kim,
VP of Strategy & Growth, Affinidi

linkedin.com/in/hongtae-kim

To learn more about other Affinidi solutions visit: www.affinidi.com