THE SECURITY SHIFT

YOUR AGENTS HAVE THE KEYS. BUT WHO SECURES THEM?

AI agents are being handed access faster than anyone can say who they are, what they may do or who answers for them. Identity and Access is now the control plane for every decision AI makes.

LONDON · SEPTEMBER 2026

changemakers.

accenture

PingIdentity

01 THE VERDICT

The agent isn't the user. Most identity systems still treat it like one.

> *"As we deploy agentic technology, we need a clear identity plane across the complete agentic lifecycle. What agents are out there, what do we want them to do and for how long?"*
>
> Jonny Hay, Strategic Manager, Ping Identity

What is clear is how little many organisations know about their own estate. Those who have invested in finding out told the same story: far more vulnerabilities and identities than expected, many surfacing only once agents arrived. For one guest, a huge share of what could be exploited sat in credentials: identity and access, pure and simple.

Agents haven't invented these identity problems. They have shone a light on them, and on how much now rests on access. Every organisation should be able to answer three questions about every identity in its estate: who is it, what can it access and what shouldn't it access?

The group came to a clear conclusion. Those answers can't be left to the agents. Decisions about access and security belong to the people accountable for them, and that responsibility grows with every agent you deploy.

BY THE NUMBERS

Nearly 2/3
of organisations identify security and risk concerns as the biggest barrier to scaling Agentic AI. (McKinsey 2026)

100:1
is the ratio nonhuman identities outnumber human ones in many organisations. (Forbes 2026)

68%
of organisations cannot reliably distinguish AI agent activity from human activity within their systems. (Forbes 2026)

Only 3%
of UK organisations are fully ready for advanced, agentic AI. (Accenture 2026)

TAKEAWAY

The exploit path already runs through identity. Agents don't open a new front, they widen the one attackers already prefer.

IN THE ROOM

The pressure from the C-suite is adopt it yesterday. 2026 is a year too late.

We're moving so fast we forget the agent isn't the user. It can go rogue. It needs guardrails and ownership.

02 WHERE IT BREAKS DOWN

Five places agents are outrunning your controls

1 The agent isn’t the user

Identity was built to answer one question about a person: who is the who? An agent is not the person who deployed it. It changes second by second and, when it can't finish a task, it may simply ask another agent to. One security architect described hundreds of agents piling onto the same problem until it looks like a denial of a service attack you launched on yourself.

2 The loop

In the loop, on the loop, at the loop, above the loop. The room laughed at the vocabulary but the point underneath was serious. As one IAM leader put it, the whole reason for a human in the loop is accountability. Move the human and you have to move the accountability with them, or it quietly disappears.

3 Old governance on new technology

One image stuck for the rest of the night: take the nozzle off a liquid soap dispenser and push it into a bar of soap. That is what bolting AI onto old infrastructure, or old governance onto AI, looks like. You must have the fundementals of governance and security in place ahead of agent integration.

4 Maturity nobody can measure

Asked if he was comfortably in control of his agents, one guest gave the most honest answer of the night: they have maturity models, but aren't sure which ones to use. Another challenged the premise. How can you score maturity when the technology itself hasn't matured? The bar moves every few months.

5 Trust you can’t observe

The challenge is not simply whether to build or buy, but how to establish trust when parts of the system sit outside your direct view. Buying offers speed and scale, while building offers greater control. Either way, organisations need enough observability across identity, behaviour and resilience to understand where trust is placed and how it is verified.

THE PATTERN

Not one of these is a tooling gap. Every one is a question of ownership: who the agent is, who it answers to and who decided it could act.

The technology is moving faster than the operating model around it.

So the fix starts with people and policy, not neccessarily product.

Make identity the control plane for AI.

There was no argument to slow down. The argument was for knowing, at every moment, what each agent is, what it may touch and which human answers for it. Get that right and autonomy becomes something you grant with confidence, not something you hope goes well.



WHAT CHANGES

Same identity disciplines, applied to a new kind of worker. Discovery, ownership and policy become the load-bearing controls.

Govern first, even if the answer is no. One security architect rejected the idea that governance can follow deployment. Start from zero trust and make each permission earn its place, including the conversation about what happens when the agent gets it wrong.

Find your agents before you secure them. Many leaders will discover far more people using AI than they think. You can't draw a perimeter around agents you don't know exist, so discovery comes first.

Give every agent a line manager. One leader said HR now wants every agent registered against a named human. Another sketched the structure: worker agents, a supervising agent reading their logs and, at the top, a person who owns the department.

Let autonomy be earned. Throttle decision making by track record. When an agent has done a task a hundred times without error, let it take three steps alone and check in on the fourth. Autonomy becomes a dial, not a switch.

Write the perimeter in policy. Fine grained policy should define what an agent may decide, refined as it proves itself. Pair it with an independent agent watching the logs, and never accept a model's own confidence score as proof. One guest noted this as the confident liar problem.

> *"Every AI agent needs an identity, an owner and a limit. Many organisations are deploying the first element before they've decided on the other two."*
>
> Ian Finn, Senior Product & Transformation Executive

04 MEASURE WHAT'S VALUABLE

Stop counting adoption. Start counting accountability.

Most AI dashboards were built to prove momentum to the board. One guest summed up the pressure as a target of 70% adoption, whatever that means. The metrics that matter now show whether the organisation still knows who is acting on its behalf.

EASY TO MEASURE - TODAY	VALUABLE TO MEASURE - INSTEAD
Contract cycle time to signature	**Revenue retained through renewal**
Templates in the CLM library	**Cash conversion from contract terms to invoice**
Pilot usage and satisfaction scores	**The business KPI defined before the pilot began**
Legal review turnaround time	**Time from missed obligation to detection**
Contracts stored in the system	**Contracts connected to delivery, billing and CRM**

On cost: The table split cleanly. A governance lead warned that nobody was talking about cost, and that today's sign-offs will need a business case in two years' time. Another leader saw the opposite, with AI testing hundreds of applications for a fraction of the old price. Both can be true. Savings in testing won't pay for the controls agents need.

On build versus buy: Building is suddenly cheap. Owning what you build is not. Resilience, identity and keeping pace with the market all land on the team that built it. The line one leader drew was simple: build where you gain an edge, buy where someone already does the job well.

> *"All organisations need to understand who is ultimately accountable for the agents. We can't let them run rogue & leave you vulnerable."*
>
> Roundtable Participant

On sovereignty: Another security leader expects the pendulum to swing back. Seamless cloud interoperability was the buzzword of 2025. As sovereignty climbs the political agenda, the idea that on premise is dead is starting to look rather dead itself.

On accountability: When an agent's decision has to be defended in court, "the agent decided" won't be an answer. You can't hold a bot accountable. Somebody has to own it.

Final thought: AI hasn't made identity harder to understand. It has made it impossible to ignore. The organisations that scale agents safely will be the ones that can say, at any moment, what each agent is, what it can access and which person answers for it

CHANGE MAKERS CLUB

Change Makers is a global community of senior executives driving long-lasting change within large, complex organisations. Members bench-test their thinking, share peer review, and discuss the opportunities and challenges of leading through transformation.

CHANGEMAKERSCLUB.COM

The insights in this report were drawn from a private roundtable of senior technology, security and identity executives, held under the Chatham House Rule and facilitated by James Harris, Founder of Change Makers.

HOSTED & FACILITATED BY

James Harris

Founder, Change Makers Club

In partnership with Ping Identity & Accenture | London, September 2026

CONTRIBUTORS

Ian Finn - Senior Product & Transformation Executive
Roy Corneloues - Security Architect, Customer Identity & Zero Trust - Hargreaves Lansdown
Conrad Lennard - Founder - Bubbling Spring
Vikas Singh - Head of Enterprise Architecture - CBRE
Mark Johnson - Head of Enterprise Wide Vulnerability Management - BNY
Victor Asanmi - Senior IAM Analyst - Experian
Ian Carpenter - Head of Data & AI Governance - Heathrow
Vitalii Chetvertakov - Senior AI Security Automation Engineer - WTW
Jonny Hay - Strategic Manager - Ping Identity
Rodrigue Perret - Manager, Security Consulting - Accenture

References

Generating Impact: Harnessing frontier AI capabilities to unlock frontline productivity and growth in the UK (Accenture 2026)

Identity As Infrastructure: Why The Future Enterprise Runs On Trust Graphs (Forbes 2026)

Cybersecurity in an Age of Disruptive Innovation (Forbes 2026)

AI Agents Aren't Creating Security Problems. They're Revealing Them. (Forbes 2026)

State of AI trust in 2026: Shifting to the agentic era (McKinsey 2026)

changemakers.

accenture

PingIdentity