changemakers

Know Your Agent

INDUSTRY REPORT

2026

PingIdentity

THE DISCUSSION TOPIC

Trust must be engineered

"Trust in agentic AI depends heavily on how the system is architected. Poorly designed systems without guardrails simply aren't trustworthy." Anon.

Trust in agentic AI does not come from the technology itself, it comes from the architecture, guardrails, and identity controls organisations build around it. Trust in AI is not a binary question of whether the technology works. Instead, it is determined by how systems are designed, governed and monitored. Well-designed system, with layered monitoring, task-based permissions, and tightly governed identities, can significantly increase confidence in autonomous systems. This shift requires organisations to treat AI agents not simply as tools, but as digital actors within enterprise systems. Just like employees, these agents require clearly defined roles, permissions, and oversight.

Actionable Solutions

"If they're constantly connected and can be tricked, they can exfiltrate data much faster than a human can." Matt Berzinski, Ping Identity

Design guardrails into AI architectures from the outset: Trust must be embedded through system design. Implement layered controls, verification mechanisms and monitoring to supervise autonomous systems.

Adopt task-based permissions and least-privilege access: AI agents should only be able to access systems and data required for the specific task they are performing.

Introduce monitoring and audit mechanisms: Organisations should track actions taken by AI agents and maintain clear audit trails to ensure they behave as expected.

Questions to Consider:

- *How confident are you that your AI systems operate within clearly defined guardrails and permissions?*
- *Are you treating AI agents as governed digital identities, or simply as tools running within your systems?*

Internal agents come first

"Some individuals trust AI quite highly, but organisations are far less comfortable, especially in regulated environments." Anon

Most organisations are starting their agentic AI journeys inside the business rather than deploying it directly to customers. Internal use cases such as HR services, knowledge assistants and operational support are seen as lower risk environments in which to experiment with AI. By contrast, external use cases raise far greater concerns. When AI interacts directly with customers, patients or financial users, the stakes become significantly higher. Reputational damage, regulatory breaches, and financial liability all become potential consequences if agents behave unpredictably. This internal-first approach allows organisations to build experience and confidence with agents before exposing them to external environments.

Actionable Solutions

"When it comes to internal usage, trust is fairly high. But when it comes to selling it to customers, most are not there yet." Abhishek Mandal, Vodafone Group

Use internal operations as testing grounds for agentic deployment: Deploy AI in controlled environments first, allowing organisations to develop governance models before scaling externally.

Define risk thresholds for external AI use cases: Not all AI applications carry equal risk. Organisations should clearly define where human oversight remains essential.

Gradually expand AI exposure through staged adoption: Start with internal augmentation, then move to decision support, before eventually exploring autonomous external use cases.

Questions to Consider:

- *Where is your organisation most comfortable deploying agentic AI today?*
- *What conditions would need to change before you trust AI to interact directly with your customers?*

Identity becomes the foundation of trust

"As AI agents begin interacting with other systems, organisations must confirm what that agent represents and whether it can be trusted." Anon

A significant challenge is how organisations verify the identity and intent of AI agents. As these systems begin communicating with other systems, customers and even other organisations, traditional identity models designed for human users become insufficient. There is a growing importance on confirming what an AI agent represents, what permissions it has, and whether it should be trusted to perform actions on behalf of an organisation or individual. Without robust identity verification frameworks, organisations risk creating complex ecosystems of interacting AI systems with limited transparency or accountability. Identity verification becomes essential to ensure organisations know who, or what, they are interacting with.

Actionable Solutions

Trust is incredibly hard to build, but so quickly and easily broken" Abhishek Mandal, Vodafone Group

Extend identity frameworks to cover AI agents: Treat agents as identities within enterprise security frameworks, ensuring their actions can be authenticated and traced.

Verify the identity of external AI systems: Organisations should ensure they understand and verify the AI systems operated by partners and suppliers.

Create clear identity-to-permission mapping: Agents should only perform actions explicitly linked to the identities they represent.

Questions to Consider:

- *How will you verify the identity of AI agents interacting with your systems?*
- *Right now, can you confidently trace every automated action back to a verified identity?*

Data maturity can build trust

"You can only trust your AI as much as the quality of your data." Anon

Organisations can only trust AI as much as they trust the quality and governance of their underlying data. Our executives repeatedly highlighted that the effectiveness of AI systems is directly tied to the quality of organisational data. Many companies are eager to deploy AI capabilities but have not yet built the foundational data governance required to support them. Poor data quality can lead to inconsistent outputs, biased recommendations and unreliable decision-making. In highly regulated industries such as banking or pharmaceuticals, these risks become particularly significant. Ultimately, organisations that invest in data maturity will be far better positioned to trust AI systems acting on their behalf.

Actionable Solutions

"The important thing is not removing bias, it's detecting that bias exists in the first place." Anon

Strengthen data governance before scaling AI: Ensure organisational data is structured, consistent and governed before deploying AI systems at scale.

Develop mechanisms to detect bias early: Regularly test AI outputs to identify patterns of bias before they impact decision-making.

Improve organisational data literacy: Ensure teams understand how AI systems rely on data quality to function effectively.

Questions to Consider:

- *Is your organisation's data foundation strong enough to support autonomous AI systems?*
- *How confident are you that your AI systems are producing unbiased and consistent outputs?*

The question of accountability

"Organisations may operate multiple interacting agents , yet no clear accountability when something goes wrong." Cigdem Cirik Arslan, GSK

As AI agents interact across organisations and systems, determining responsibility for their actions becomes increasingly complex. Executives question how organisations will assign responsibility when agents make mistakes or interact with other automated systems. In complex digital ecosystems, multiple agents may interact across organisational boundaries, making it difficult to determine where errors originate. This creates both legal and operational challenges, particularly in highly regulated industries. Executives stressed that while AI systems may operate autonomously, organisations must still ensure that clear accountability frameworks exist. Without this, businesses risk deploying systems where responsibility cannot be easily traced when failures occur.

Actionable Solutions

"The real challenge will be determining responsibility when multiple AI systems interact across organisations." Anon

Maintain human accountability for AI decisions: Even when AI systems operate autonomously, organisations must assign clear human responsibility for outcomes.

Implement traceability across AI systems: Ensure you can trace decisions made by AI systems through logs and monitoring frameworks. Bring in an identity framework that can work across these multiple complex systems.

Define governance models for cross-organisation AI interactions: As AI agents begin interacting across companies, you must establish shared accountability frameworks.

Questions to Consider:

- *Who is accountable when an autonomous AI agent makes a mistake?*
- *Can your organisation trace how an AI system arrived at a decision?*

Conclusion

The discussion made one point clear: as organisations move from automation to autonomous AI agents, trust cannot be assumed, it must be engineered. Agentic AI introduces a new layer of complexity to digital ecosystems, where non-human identities can access systems, interact with other agents, and make decisions without direct human intervention. While the opportunity for efficiency and innovation is substantial, executives agreed that the risks increase at the same pace. In this environment, organisations must rethink how they verify intent, control permissions and monitor activity, ensuring that every AI agent operating within their systems can be trusted to act within clearly defined boundaries.

A recurring theme throughout the roundtable was that existing identity and access models were largely designed for human users. As AI agents become active participants in enterprise workflows, these models must evolve. Organisations will need far greater visibility into what autonomous agents represent, what permissions they hold, and how their actions can be traced back to a responsible identity. Without robust identity and access management frameworks, businesses risk creating ecosystems where autonomous agents operate with unclear ownership, excessive permissions, or insufficient oversight, introducing security, compliance and reputational risks.

Ultimately, the future of agentic AI will depend not only on the power of the technology, but on the discipline with which organisations manage it. Strong identity governance, least-privilege access, and clear accountability structures will be essential to building confidence in autonomous systems. As AI agents begin interacting with customers, partners and other organisations, the ability to verify and govern these digital identities will become foundational. Those organisations that succeed will be the ones that combine innovation with rigorous identity and access management, ensuring that the autonomy of AI is matched by the security and trust needed to support it.

KNOW YOUR AGENT

Change Makers is a global community of senior executives who aspire to drive long-lasting change within their large and complex organisation. As a community we discuss the opportunities and challenges ahead, bench test our thinking and hear peer review. On a regular basis we meet for confidential meetings and produce industry reports. Learn more about the Club, access content and apply to join up-coming events via our website.

www.changemakersclub.com

AUTHORS

James Harris
Club Founder

George Whittington
Club Director

Ping Identity believe in making digital experiences both secure and seamless for all users, without compromise. That's digital freedom. They let enterprises combine their best-in-class identity solutions with third-party services they already use to remove passwords, prevent fraud, support Zero Trust, or anything in between. This is all made possible using their simple drag-and-drop canvas. That's why more than half of the Fortune 100 choose Ping Identity to protect digital interactions from their users while making experiences frictionless.

Thank you to our thought leaders for the session:

Matt Berzinski, EMEA Field CTO, Ping Identity

Jonny Hay, Strategic Account Manager

Tom Mendlinger, Strategic Account Executive