CISCO X CHANGE MAKERS

WHO SIGNS OFF THE ROBOT?

The Intersection of AI Governance and Infrastructure in Financial Services

AMSTERDAM | JUNE 2026

changemakers.

cisco

01 THE STATE OF PLAY

AI has stopped being a technology problem

Held under Chatham House rules, the conversation opened from a shared premise: AI has crossed from a purely technical concern into a paramount, board-level issue, one that now bridges the CISO, first-line risk and GRC functions.

A recurring frustration surfaced early: there is still no harmonised way for institutions to tier technology risk. Two organisations building near-identical AI use cases can land on entirely different risk classifications, with no shared reference point to reconcile them.

That ambiguity has real consequences for who is willing to move. For large, established European banks, the downside is asymmetric: a neobank might absorb a modest fine for getting AI wrong, but an established institution risks billion-dollar penalties and lasting reputational damage, hence the hesitation to scale AI past proof-of-concept, even as the pace of the technology itself is, in more than one participant's words, “insane.”

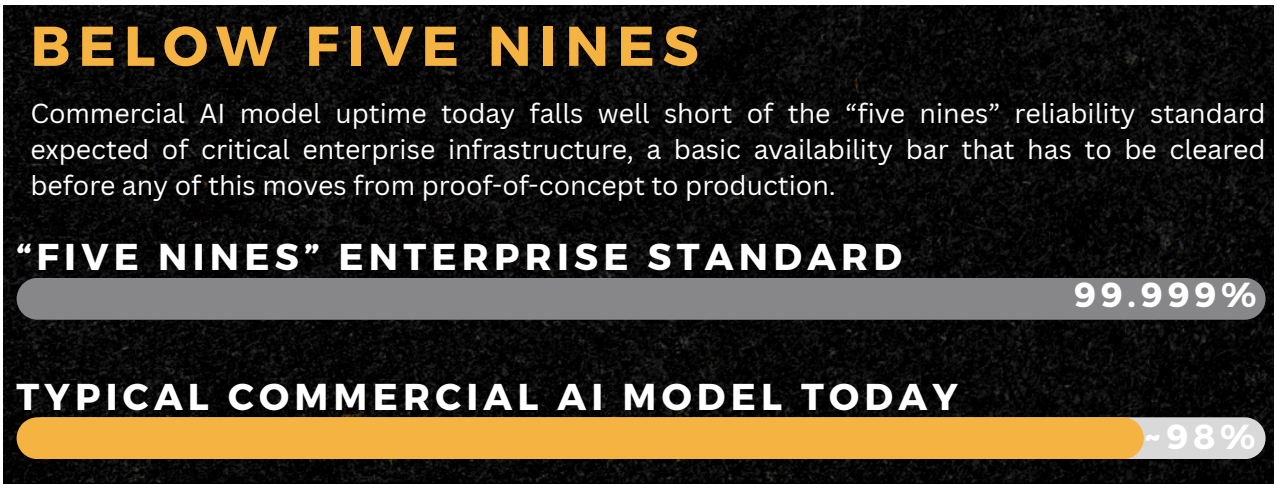

> *“For a large, established bank, this isn't a fine, it's a billion-dollar penalty and your reputation gone. That's why nobody wants to be first out of proof-of-concept.”*
> Roundtable Participant

10-SECOND TAKEAWAY

AI is now a board-level governance issue, not just a technology one, but the industry still has no shared, harmonised way to tier its risk. Until that exists, the biggest, most capable institutions have the most to lose by moving first.

02 WHERE IT BREAKS DOWN

Five patterns undermining AI adoption

Strip away the hype and the same handful of structural problems came up again and again, all of them rooted in governance, infrastructure or trust rather than the technology itself.

1 The Human-in-the-Loop Trap

Regulators still demand a human sign-off on AI-driven decisions like transaction monitoring. It feels safer, but participants agreed it isn't sustainable given daily transaction volumes, and human review is often the actual point of failure, prone to emotional bias and error that the model itself wouldn't make.

2 Silos Across the Business

Centralised platform teams want business domains to experiment freely, but every reusable AI component still has to pass through rigorous audits and second-line risk assessment before it can scale, slowing standardisation to a crawl even where the appetite to move fast exists.

3 Platform Fragmentation

One institution described running Microsoft's platform for its workforce, Google's Agentic Enterprise for processes like mortgages, and GitHub for 18,000 developers. Each has real strengths, but nobody has resolved who is responsible for controlling deterministic agent behaviour across all three.

4 The Power Wall

A single AI server rack is expected to draw 100 kilowatts within three years, roughly the power draw of a small car. Grid capacity, not compute, is now the binding constraint: new connection permits can take up to 14 years to secure in some regions.

5 Invisible Agents

Most organisations cannot yet produce a full inventory of the AI agents already running inside their business, let alone assign each one a named human owner or scope its access. That makes least-privilege, zero-trust access aspirational rather than actual, for now.

03 WHAT'S WORKING

Five moves the room is actually taking

None of these are silver bullets, but they're what separated the organisations making real progress from the ones still stuck debating strategy.

1 Build a Real Centre of Excellence

Not just a governance layer, a genuinely multidisciplinary team, where AI engineers become more business-fluent and auditors develop real technical understanding of the models they're reviewing. That shared language is what actually shortens the audit bottleneck over time.

---

2 Treat Data as the Foundation, Not an Afterthought

The institutions furthest ahead had already done the unglamorous work of identifying and classifying their most critical data assets, rather than waiting for AI itself to solve indexing and classification, which several participants expect within two years, but can't wait for.

---

3 Bring Security in at Design, Not at Launch

Algorithmic pen-testing of commercial models routinely surfaces fundamental security issues, but security teams are still being brought in too late in the development cycle, forcing costly rework that earlier involvement would have avoided entirely.

---

4 Inventory Every Agent, Then Apply Zero-Trust

You can't secure what you can't see. The starting point is a full inventory of every agent already running in the business, followed by least-privilege access and a named human owner for each one, the same standard already applied to human identities.

---

5 Collaborate on Shared Standards

The industry has historically moved at a similar pace when it works together. Continued dialogue between institutions and technology partners on common standards for risk, governance and infrastructure was seen as the clearest way to avoid a two-tier competitive landscape.

FINAL THOUGHT

The bottleneck isn't the model. It's everything around it.

Nothing in this conversation suggested the technology is the limiting factor. What's holding institutions back is regulation with no shared risk taxonomy, infrastructure that physically can't keep pace with agentic AI's power demands, and a security paradigm that hasn't caught up to a world of autonomous, largely invisible agents. The banks moving fastest aren't necessarily the most advanced technically — they're the ones treating governance, data and infrastructure as part of the AI build, not obstacles to it.

That's also where the competitive risk sits. A small group of the largest banks are already scaling agentic AI well ahead of the rest of the market, echoing the same "too small to compete" dynamic the industry saw in the years after the last financial crisis. Geopolitical pressure is now pushing European institutions to build sovereign AI capability of their own. The consensus in the room was that no single institution solves this alone: continued, practical collaboration between banks and technology partners on shared standards is what keeps the playing field level.

> *"We cannot even identify all the agents running in our business, let alone assign them an owner. That has to be the starting point, not the end goal."*
>
> Consensus View, Cisco Amsterdam Roundtable

Want to read more?

Read more of our reports on the Change Makers website and see what future events we have coming up.

Further Insights →

CHANGE MAKERS

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 financial services executives and industry leaders, facilitated by James Harris, Founder of Change Makers.

FACILITATED BY

James Harris

Founder, Change Makers Club

In partnership with Cisco, Amsterdam June 2026

CONTRIBUTORS

Nassim Mehrshid - Head of AI - Finance & Risk at ABN AMRO
Bart Jansen - Agentic AI Transformation Director at ING
Shan Amin - AI Transformation Lead at NN
Sashidhar Mikkilinenili - Senior DevOps Engineer at ING
Raj Dash - Director of Data Management - Financial Markets at ING
Tijana Nikolic - AI Technologies Product Owner at ABN AMRO
Joram Frieser - Lead Data Governance Officer at BNG
Moustafa ELZeny - Platform Engineering Lead at Rabobank
Christian Vermorken - Domain Architect at Consistent Solutions
Shruti Naldiga - Technical Product Owner at ABN AMRO
Steffanie van der Maas - Chief Auditor at Rabobank
Al Slamecka - FSI Global Cloud & AI Infrastructure at Cisco
Koen van den Berg - AI Factory Specialist EMEA at NVIDIA
Thomas Vrancken - Senior ML Lead at ML6
Jeff Denworth - Co-Founder at VAST Data
Jan Heijdra - Field CTO Security & AI at Cisco
Bas Smit - Enterprise Sales Manager at Cisco
Rens Verheijde - Enterprise Strategy Account Executive at Cisco

changemakers.

cisco