Trailblazer

aws FOR SOFTWARE AND TECHNOLOGY

GenAI Trailblazers

Future-Minded. Value-Driven.

Roundtable three I Debriefing paper
October 2025

GenAI has a trust problem. The question now is: how do we solve it? Trailblazers from Cato Networks, Genesys, Intercom, Mimecast, and Monday.com share their insights.

A Meet the Boss community
in partnership with AWS

Meet the Boss

gds

Future-Minded. Value-Driven.

Contents

Introduction.

Generative AI has reached a pivotal moment.

While adoption continues to accelerate, trust has not kept pace. [A recent survey by Deloitte of 30,000 Europeans]() found that deepfakes, misinformation, data privacy and inaccurate results remain the top public concerns. Perhaps that explains why GenAI use for personal activities (47%) still far outstrips its use for work (23%).

At the same time, the threat landscape is evolving fast. Cyberattacks targeting software applications have increased by 300%, with adversaries already utilising AI to automate vulnerability discovery and craft highly realistic phishing campaigns.

For ISVs, cybersecurity is becoming more than a defensive necessity – it's a strategic and competitive differentiator.

This roundtable brought together members of the Trailblazer Community to explore how trust can be built into the very fabric of AI systems. How do you govern what you can't fully predict? What defines "responsible" AI in practice? And how can organisations deliver innovation without compromising integrity?

GenAI has a trust problem. The question now is: how do we solve it?

Roundtable three

Avidan Avraham,
Security Researcher,
Cato Networks

Azam Banaras,
Director, Global IT End User Services & AV, Genesys

Terri Parker,
Senior Product Manager,
Mimecast

Deven Patel,
Director of Product Analytics,
Intercom

Chen Salomon,
Director of Engineering,
Monday.com

Ruslan Semenov,
Director of Engineering,
Monday.com

"Technology is a useful servant but a dangerous master."

Christian Lous Lange, Historian and Teacher

Chapter 1

Governance and the human factor

Trust, the group agreed, isn't built on technology alone. It relies on governance, design, and human judgement working together – and on a willingness to see compliance not as a barrier to progress, but as the system that ensures progress is sustainable.

> **"We can't expect users to care enough about data privacy. We have to care on their behalf."**
>
> **Deven Patel,** Intercom

Deven Patel of Intercom highlighted the paradox at the core of AI adoption: users claim to value privacy, yet their behaviour indicates that convenience often prevails. "People still click 'accept all'," he said. "So it's on us to care on their behalf." For Deven, governance is more about alignment than strict procedures – ensuring clear accountability for where data is stored, how it's used, and who is responsible if something goes wrong.

Terri Parker from Mimecast agreed, explaining that consumer education can only go so far. "You can't train everyone to be an expert in data hygiene," she said. "That's why governance and automation have to carry the load." Mimecast's strategy is to embed guardrails directly into tools and workflows, so security becomes part of the experience rather than an obstacle to it.

Chen Salomon from Monday.com added that regulation has become the main factor driving AI maturity. "Legislation sets the minimum standard and compels you to formalise processes that might otherwise remain informal," he said. However, he warned against excessive governance: too many overlapping frameworks could impede experimentation. The challenge is proportionate, ensuring controls align with the level of risk rather than unnecessarily slowing progress.

From an operational perspective, **Azam Banaras** at Genesys highlighted a growing governance gap between Europe and the US. European organisations tend to prioritise compliance from the outset, he noted, while many American firms retrofit governance later on. For global ISVs, this necessitates designing policy frameworks versatile enough to satisfy both approaches without significantly increasing overhead.

Avidan Avraham of Cato Networks observed the same trend in security. “We’re running faster than the regulators,” he said. “So we have to self-govern.” He highlighted Cato’s focus on early risk classification, incident response drills, and AI-assisted anomaly detection – practices that anticipate future compliance requirements rather than waiting for them.

Throughout the discussion, a consensus emerged that trust must be engineered, not assumed. Policies, training, automation, and culture all work in concert. Terri Parker summed it up: “Education alone won’t stop a breach. Governance and guardrails have to do the heavy lifting.”

Governance is no longer about saying “no” to innovation; it’s about ensuring organisations can keep saying “yes” safely and repeatedly.

Trust and Agentic AI

Ruslan Semenov from Monday.com raised the question of how much organisations can trust autonomous agents to act safely. He stated that the more tools an agent has, the more unpredictable its behaviour becomes. “You’re giving it access, not just instructions,” he warned. Therefore, transparency and monitoring must evolve alongside capability. For Ruslan, responsible use of agentic systems begins with visibility, such as clear logs of every action and an auditable chain of decisions. Visibility enables teams to spot when autonomy crosses a line.

> **“Education alone won’t stop a breach. Governance and guardrails have to do the heavy lifting.”**
>
> **Terri Parker,** Mimecast

Can you fight fire with fire?

Should GenAI be used to defend against AI-powered attacks? **Avidan Avraham** argued yes: “Attackers automate, so defenders have to automate too.” AI can detect anomalies faster and on a larger scale than humans ever could. **Terri Parker** agreed in part but warned that defensive AI must still operate within strict governance: “Otherwise you’re just creating a new layer of risk.” **Deven Patel** added that explainability is critical – if a system flags a threat, humans need to understand why. **Ruslan Semenov** concluded that “zero trust” principles still apply, even to AI itself: verify, test, and audit continuously.

Chapter 2

Seven challenges on the road to trust

If governance lays the groundwork, the next step is to build on it and turn principles into practice. When our group considered the real-world challenges of trustworthy AI, seven persistent issues arose: proof of value, responsible adoption, bias, transparency, trust, accuracy, and risk management.

Chen Salomon began with proof of value. Too many AI projects, he said, scale before they demonstrate sustained benefit. "A pilot that saves time once isn't proof that it will keep doing so at scale." Monday.com assesses long-term value through operational metrics such as the friction removed from everyday work.

Azam Banaras highlighted the global nature of responsibility. "Governance can't stop at borders," he said. "Customers don't care whether you're bound by EU or US law; they just want to trust you." That involves harmonising privacy, transparency, and explainability standards across regions.

Terri Parker emphasised bias and fairness. Every training set reflects the human decisions that inform it. "You can't remove bias completely," she said, "but you can make it visible." Mimecast now conducts bias audits alongside security assessments, making them a shared responsibility between data, design, and compliance teams.

Deven Patel emphasised transparency. As systems grow more complex, trust hinges on the ability to explain how decisions are made. "At some point, you have to open the box," he said. "The more context you can show, the more people will believe the outcome." He maintained that transparency is not merely a feature but an architecture that integrates engineering, product, and legal teams from the outset.

Avidan Avraham revisited the issue of accuracy. "The more tools you give an agent, the more confused it becomes," he said. "Trust starts with focus." He warned that organisations pursuing a broad range of capabilities often compromise reliability, and that clear task boundaries are essential for maintaining verifiable systems.

Ruslan Semenov concluded with a focus on data quality and risk prioritisation. "Your AI is only as good as your data," he said. "Noise costs money and credibility." Monday.com now applies the same maturity model to data governance as to product development: define, test, iterate, secure.

> **"A pilot that saves time once isn't proof that it will keep doing so at scale."**
>
> **Chen Salomon,** Monday.com

The seven challenges of trustworthy AI

What Trailblazer leaders say must come next…

CHALLENGE		KEY INSIGHT
Proof of value	1	Scale only what delivers lasting impact
Responsible adoption	2	Governance has no borders
Bias and fairness	3	Make bias visible, not invisible
Transparency	4	Explainability builds belief
Trust	5	Accountability across functions
Accuracy	6	Trust starts with focus
Risk focus	7	Noise costs money and credibility

“Your AI is only as good as your data.”

Ruslan Semenov, Monday.com

Conclusions

For these Trailblazers, trust is an ecosystem, says community manager **Adam Burns.**

Our roundtable concluded with a shared belief: building trustworthy AI is more about culture than technology. Governance sets boundaries, but it is human integrity that maintains honest innovation.

For ISVs, this shift is redefining competitive advantage. Security, transparency, and compliance are now central to customer value – not just supporting elements but key differentiators. And trust, once taken for granted, is becoming a product in its own right.

And a note of quiet confidence: the organisations that succeed will be those that prioritise safety, communicate clearly, and keep curiosity alive within strong governance. That's what makes AI not just powerful but deserving of trust.

Trailblazer

Future-Minded. Value-Driven.

gds | aws FOR SOFTWARE AND TECHNOLOGY