

AI Agents & Platform Engineering: Efficiency Boost or New Source of Trouble?



Hasith Kalpage

Director, Platform
Engineering & CISO



Idit Levine

Founder & CEO



Vincent Caldeira

CTO APAC



Sara Qasmi

Platform Engineer



Carlos Santana

Solutions Architect
for EKS





AGENDA

- Use Cases
- Challenges
- Trust
- Moving forward



AGENTIC USE CASES

- Automated incident response & root cause analysis
- Infrastructure provisioning & self-healing systems
- CI/CD pipeline optimization & deployment automation
- Security scanning, compliance checks & policy enforcement
- Developer self-service portals powered by AI agents



TOP CHALLENGES

- Observability gaps — understanding what agents actually do
- Blast radius control — limiting damage from autonomous actions
- Hallucination risk in infrastructure-as-code generation
- Governance & audit trails for agent-driven changes
- Skill gap — teams need new expertise to manage AI workflows



LEVELS OF TRUST

- Level 0: Human-in-the-loop — agent suggests, human approves
- Level 1: Guardrailed autonomy — agents act within strict policies
- Level 2: Supervised autonomy — agents execute, humans review
- Level 3: Full autonomy — agents own end-to-end workflows
- Key question: How do you decide what level fits your org?



HOW TO GET STARTED

- Start small — pick one well-scoped, low-risk workflow
- Build guardrails first: RBAC, sandboxing, rollback mechanisms
- Invest in observability before scaling agent adoption
- Create feedback loops — measure agent effectiveness continuously
- Foster a culture of experimentation with safe failure modes