

Operationalising AI: Observability and Business Readiness

AI is transforming how engineering teams build, deploy, and scale, but real-world adoption brings complex challenges around governance, cost, code quality, and culture.

We brought together a panel of Engineering leaders to share what's actually working when operationalising AI, including observability and business readiness. Check out the key takeaways below:

SESSION LED BY:



Robert Heywood

AI Engineering Leader



BUILDING AN AI-READY ORGANISATION

AI Ambassadors

- Establish internal AI “champions” across different disciplines, especially non-technical, to trial, test, and share best practice.
- Form cross-functional AI communities to create momentum and visibility.
- Embed AI into objectives, don’t leave it “off the side of the desk.”
- Regularly showcase success stories to normalise adoption and encourage curiosity.

Structured Governance

- Creating a dedicated AI procurement board helps speed up safe approval of new tools.
- Form an AI SWAT Team (cross-functional specialists) to assess, pilot, and standardise.

Cultural Change & Upskilling

- Treat AI enablement as a business-wide skill. Allocate dedicated time for experimentation.
- Leadership should visibly model AI use, to encourage wider adoption.
- Address engineer fears around “replacement” by reframing AI as not a threat.
- Start with safe, valuable use cases (e.g., documentation, incident analysis, or research).

IMPROVING RELIABILITY OF AI

Observability & Measurement

- Track usage, edit rates, and return usage to gauge value and trust.
- A/B test AI features for measurable business impact.
- Use “happiness thresholds”, AI will never be 100% accurate.
- Educate stakeholders on non-deterministic systems and probabilistic outputs.

CHANGING DEVELOPMENT CYCLES

Traditional: Wireframe → Code → Test → Release

AI-first: Prompt → Prototype → User Test → Refine → Deploy

- Engineers now spend more time writing natural language than code — prompting is the new skill.
- Design and engineering are converging; collaborative iteration is now the norm.

BUILD VS. BUY DECISIONS

- Build for internal productivity; buy for scale.
- Assess inference cost before deploying to customers.
- Avoid tool FOMO, only adopt if output is meaningfully better, not just by 10%.

AI SECURITY

- Guard against prompt injection and data leakage.
- Sandbox new tools and limit free-text inputs.
- Define clear rules on human review of AI outputs.

ETHICAL GUARDRAILS

- Transparency:** Clearly identify when users are interacting with AI rather than a human, particularly important for voice AI and customer support scenarios.
- Empathy vs. Manipulation:** AI that sounds or behaves human-like (e.g., ElevenLabs voice tools) can increase engagement but raises ethical questions. Keep tone empathetic, not emotionally manipulative.
- Brand Protection:** Ensure AI systems don’t generate content or responses that could harm reputation. Train models on brand-safe data and monitor regularly for drift or off-brand behaviour.

