



OEUK HSE CONFERENCE · 25 FEBRUARY 2026

Enabling Successful Adoption of AI in Safety Management

A Framework for Success

P&J Live, Aberdeen



Arrash Nekonam

Chief Technology Officer, COMET

arrash.nekonam@cometanalysis.com

cometanalysis.com





INTRODUCING

Meet Jock McSafety

HSE Manager · 25 years offshore · Seen it all

Jock has managed safety across permits, risk assessments, audits, incident investigations, and three different reporting systems. He's competent, experienced, and deeply committed to keeping people safe. But right now, he's overwhelmed, buried under mountains of data, chasing overdue actions, and staring down yet another audit cycle.

10K+

Safety Records

Accumulated across multiple systems and assets

47

Overdue Actions

Slipping through the cracks under workload pressure

6wks

Until Audit

A major regulatory audit with no time to spare

"He's drowning in data, but he doesn't trust AI." If your AI adoption strategy doesn't work for Jock, it doesn't work at all.

THE CHALLENGE

Jock's Problem

The human cognitive limit in safety management

The volume of safety data has exploded, but our ability to process it hasn't scaled. Jock is managing incidents, risk registers, permits, inspections, audits, management of change, and contractor oversight - across systems that don't talk to each other. Critical connections go undetected.



Data Volume

Incidents, risk assessments, permits, audits, inspections- no one can cross-reference 15 years of records to find what connects them.



Ageing Assets

Risk profiles deepen as infrastructure ages. Barrier degradation, corrosion, SIL failures and knowledge leave with every crew change.



Workforce Pressure

Fewer people, same scope: investigations, permits, compliance, MoC, contractor oversight. Something always falls through.



Fragmented Systems

Incidents in one system, permits in another, risk registers in a third. No single view of safety across the asset.





THE TRUST BARRIER

Why Jock Says No

And why he's right to be sceptical

Jock's objections aren't irrational fears; they're the legitimate concerns of a competent professional. This is where most AI adoption fails: not because the technology doesn't work, but because the implementation doesn't earn trust.



"It'll replace us."

Fear of displacement is real and must be addressed head-on. If people believe AI threatens their role, adoption dies before it starts.



"It doesn't understand offshore."

Generic AI can't read a permit, interpret a bow-tie, or understand barrier management. Domain expertise is non-negotiable.



"Who validates the output?"

Without governance, AI in safety management is indefensible. Every output needs a competent human review.



"What about the regulator?"

Auditability and explainability aren't optional; they're required. If you can't explain it, you can't defend it.



This is where most AI adoption fails. So let me show you a framework that earns trust, pillar by pillar.

Earning Jock's Trust

Five pillars for successful AI adoption in safety management

These five pillars are sequential; you can't skip ahead. Each one builds on the last, creating the foundation of trust that safety professionals require before they'll accept AI into their workflows.



Problem-First Design

Start with the safety question



Data Governance

Trust the data before the output



Human-in-the-Loop

AI recommends, humans decide



Explainability & Audit Trail

Defensible under regulatory scrutiny



Incremental Value Proof

Small wins first, scale with confidence



 PILLAR 1

Problem-First Design

Start with Jock's problem, not the technology

✗ Wrong Approach

"We've got AI, what can we use it for?"

Technology looking for a problem. This approach fails in safety-critical contexts because it prioritises novelty over need. Without a clear safety question driving the application, you end up with expensive tools that solve no real operational challenge.

✓ Right Approach

"What safety risks are we not seeing fast enough?"

Safety question first. Then ask: can AI answer it better, faster, more reliably than current methods? Ask Jock what keeps him awake at night. Build the AI around that.

Questions that drive the right AI applications:

- Which overdue actions are linked to our highest-risk scenarios?
- Are there patterns connecting our incidents, near-misses, and permit deviations?
- Where are our barrier controls degrading faster than our inspection cycles detect?

Data Governance

Jock needs to trust the data before he'll trust the output

Garbage in, garbage out. Jock knows this instinctively. Before AI can deliver value, you must confront the state of your safety data honestly. The same hazard was classified three different ways across incidents, permits, and risk assessments. Free-text narratives that neither machine nor human can compare consistently.

Inconsistent Taxonomies

Same hazard classified differently in incidents, risk registers, and permit systems. No common language across the safety lifecycle makes cross-referencing nearly impossible.

Free-Text Chaos

Critical detail buried in narratives, inspection notes, and MoC justifications, no two people describe the same event the same way, making pattern recognition a manual nightmare.

Legacy System Silos

Incidents in one system, permits in another, risk assessments in a third, inspections in a fourth. No single view of safety performance across the asset.

- Key insight: AI forces you to confront your data quality, and that's a benefit, not a burden. The governance you build for AI improves all your safety reporting and regulatory submissions.





👤 PILLAR 3

Human-in-the-Loop

Jock stays in control, AI recommends, humans decide

In safety-critical contexts, AI is an analyst, not an authority. AI can process, pattern-match, and prioritise at scale. But the decision to act, to close an action, approve a permit, change an inspection regime, or escalate a risk must rest with a competent person who understands the operational context.



ALARP

AI helps demonstrate risk is as low as reasonably practicable, surfacing evidence humans might miss across thousands of records and multiple systems.



Safety Case

AI supports the evidence base for your safety case; it doesn't replace engineering judgement. The human expertise in interpreting context remains irreplaceable.



Competence

The reviewer must understand what AI is showing them. Training and competence frameworks must evolve to include AI literacy alongside technical safety skills.

Explainability & Audit Trail

Jock can explain it to the regulator

The Regulator Test

"Can you explain how this conclusion was reached?"

If Jock's answer is "the AI told me", that is not defensible. But if he can show source data, methodology, confidence level, and his decision with a rationale that IS defensible.

Black-box models fail this test. Explainable AI isn't a nice-to-have, it's the cost of entry.

Every AI output in safety management must answer three critical questions:

01

What data went in?

Full traceability to source records, dates, and data quality scores. No hidden inputs.

02

What logic was applied?

Transparent methodology - not a black box. Confidence levels clearly stated and understood.

03

What did the human decide?

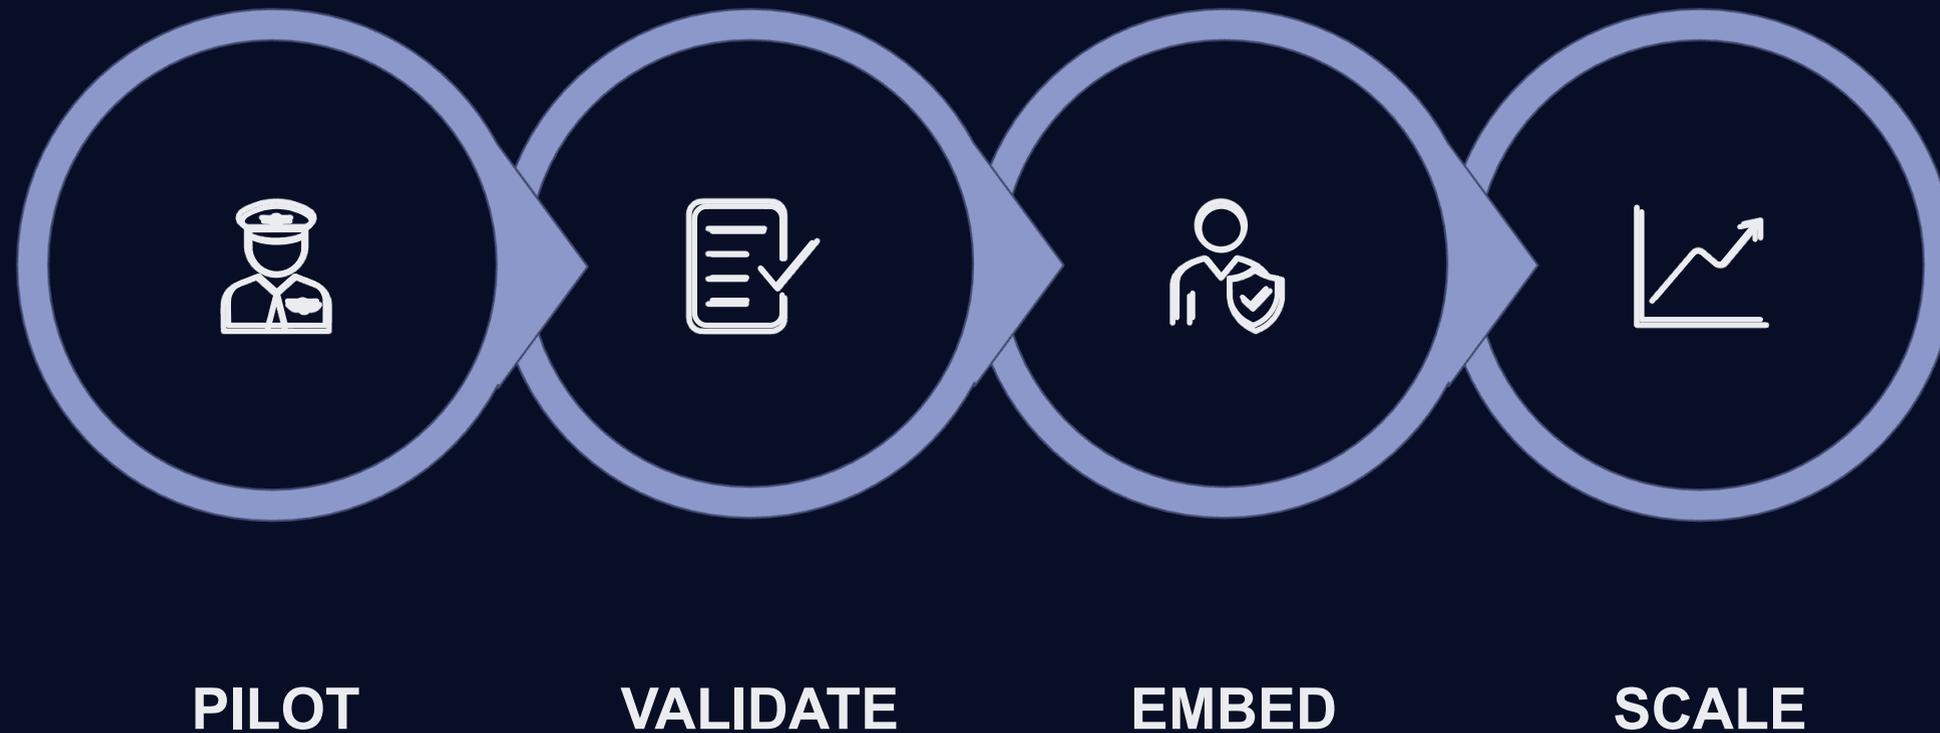
Record of acceptance, rejection, or modification with rationale documented for audit.



Incremental Value Proof

Show Jock a result he couldn't have found manually

Don't try to transform everything at once. Pick one painful, manual, high-value process. Apply AI. Show a result that changes the conversation.



The Jock Moment: "AI found 3 unlinked incidents sharing a common barrier failure and flagged 12 permits with the same risk exposure. Two years of data across three systems, connected in seconds. No human could have done that."

That's when trust begins.

THE TRANSFORMATION

Jock's Monday Morning

Not replacing Jock. Making Jock better.

✗ Before

- Hours reviewing reports and permits manually
- Risk patterns hidden across disconnected systems
- Actions overdue, barriers degrading undetected
- Reactive, always responding to the last incident
- Audit prep is panic mode

✓ After

- AI-prioritised risks and actions ready each morning
- Cross-system patterns surfaced automatically
- Barrier health monitored proactively
- Leading indicators, not lagging responses
- Audit-ready and defensible at all times

📄 The real win? Jock spends less time at his desk wrestling spreadsheets and more time on the platform talking to crews, walking the asset, managing risk where it actually happens.

BEFORE



AFTER





SUMMARY

The Framework for Success

Enabling successful adoption of AI in safety management



Problem-First Design

Safety question drives the AI



Data Governance

Trust the data before the output



Human-in-the-Loop

AI recommends, humans decide



Explainability & Audit Trail

Defensible under scrutiny



Incremental Value Proof

Small wins, then scale

Trust

Earning Jock's confidence through human-centred design

Transparency

Explainability, audit trails, and data governance

Transformation

Incremental, measurable, real-world safety improvement