

What does an AI-enabled organisation actually look like?

In June, Hancock Iron Ore Chief Technology and Information Officer Kate Flanagan published an interesting account of how AI is being used across the business.

There are plenty of individual use cases worth looking at.

AI helping anticipate equipment failures. Computer vision identifying microscopic defects in rail infrastructure and train wheels. AI estimating stockpile volumes. Models helping optimise operational parameters.

But it wasn't any one of those that caught my attention.

It was everything sitting around them.

Hancock describes AI being used across safety, maintenance, daily operations and planning.

AI resources are embedded with operational and maintenance teams.

More than 1,000 people had completed the first level of its AI literacy program at the time of publication.

There are governance and ethical guardrails.

And importantly, the technology is being connected to operational outcomes including safety, reliability, throughput and better decision-making.

That combination matters.

Because organisations can easily treat each of those things as separate pieces of an AI program.

01

Build the technology.

02

Run a change program.

03

Develop the skills.

04

Put governance around it.

05

Work out the operating
model.

06

Measure the value.

Each one makes sense.

But the more AI becomes part of how work actually gets done, the less useful I think it is to design them separately.

Take predictive maintenance.

If AI identifies an emerging equipment problem, what happens next?

-
- 01 Does it fit naturally into the way the maintenance team already works, or have we created another system people need to work around?
 - 02 What can it recommend, and what can it decide?
 - 03 When does someone need to intervene?
 - 04 What does the person receiving the recommendation need to see before they are confident enough to act?
 - 05 What evidence should the system create automatically?
 - 06 And ultimately, what changes in reliability, downtime or maintenance performance will tell us this was worth doing?

Those aren't separate workstreams around an AI implementation.

They're design inputs to the AI-enabled process itself.

That's what I find interesting about the Hancock example.

Not that they've found more places to use AI.

It's that their experience gives us a glimpse of what starts to happen when AI moves beyond isolated use cases and becomes part of the way an organisation operates.

The opportunity isn't to make all of that bigger or more complicated.

It's the opposite.

Be deliberate about the few things that matter for each use case, design them together, and then build and learn.

- Technology changes.
 - Work changes with it.
 - Roles change.
 - Decisions change.
 - Capability changes.
 - Governance changes.

And eventually, the organisation itself starts to change.

Think big. Start small.
But know what you're designing towards.



A DIFFERENT WAY TO THINK ABOUT IT

An AI-enabled organisation isn't an organisation with more AI.
It's one where AI is designed into the way the organisation operates,
alongside its people, processes and technology, to **realise value**.