

The Human Premium.

Clear thinking in complicated times.

A SELF-SCORED READINESS DIAGNOSTIC

Twelve questions before you automate anything.

Most organisations ask how to adopt AI. Very few ask first whether they should. These are the twelve questions that come before that decision — and the honest answers are rarely the comfortable ones.

Adam Clarke · The Human Premium

info@thehumanpremium.co.uk

thehumanpremium.co.uk

BEFORE YOU BEGIN

How to use this.

This takes about fifteen minutes if you answer quickly, and rather longer if you answer honestly. The second is the version worth doing.

Score each question from zero to three:

- 0 No, or I do not know.
- 1 In part, or in one corner of the organisation.
- 2 Mostly, though it rests on individuals rather than on anything written down.
- 3 Yes, and I could demonstrate it to somebody who doubted me.

Three is a high bar deliberately. If you find yourself reaching for two, take the two. The score is only useful if it is true.

The twelve questions are arranged in four groups of three, following the sequence in which these problems actually arrive: judgement, then governance, then adoption, then growth. Score each group separately as well as in total. Where the weakness clusters matters more than the number at the end.

This diagnostic is a self-report. It will locate you. It will not tell you what to do first — see page six for why that distinction matters.

GROUP ONE

Judgement.

What must remain yours.

1. Can you name, without pausing to think, the three decisions in your organisation that must never be made by a machine — and could a colleague name the same three?

Score 0 1 2 3

2. When somebody proposes an AI tool, does the conversation begin with whether the task should be automated at all, or with which tool to buy?

Score 0 1 2 3

3. Is there a named person whose refusal would halt an AI deployment, and does everybody know who that person is?

Score 0 1 2 3

Group one subtotal, out of nine:

GROUP TWO

Governance.

What you could prove if you were asked.

4. If a client, funder, or regulator asked tomorrow where AI had touched their work, could you evidence it accurately within a working day?

Score 0 1 2 3

5. Do you know which AI tools your people are already using, including the ones nobody has told you about?

Score 0 1 2 3

6. Is there a written record of who approved each system in use, and what they were told at the point they approved it?

Score 0 1 2 3

Group two subtotal, out of nine:

GROUP THREE

Adoption.

What is actually happening, rather than what was announced.

7. When a new AI tool arrives, does it go through a deliberate trial — a defined question, a set duration, and a decision at the end — or is it tried informally and then quietly kept?

Score 0 1 2 3

8. Do your people believe that raising a concern about an AI tool is safe — and can you point to an occasion when somebody did?

Score 0 1 2 3

9. When a system produces an output, does somebody with the expertise to disagree with it actually check it, or does checking mean reading it?

Score 0 1 2 3

Group three subtotal, out of nine:

GROUP FOUR

Growth.

What the gains are costing you.

10. Can you state clearly what your organisation would lose, rather than gain, from its next planned automation?

Score 0 1 2 3

11. Has the time saved by AI so far been reinvested in something you could name, or has it quietly disappeared into the working day?

Score 0 1 2 3

12. If your most capable people left next month, would the capability walking out of the door be the same capability you have been automating?

Score 0 1 2 3

Group four subtotal, out of nine:

Total, out of thirty-six:

YOUR RESULT

Where the score places you.

Four bands. None of them is a verdict, and none of them is a compliment.

0–11 Exposed.

AI is in your organisation whether or not you have approved it, and you would struggle to describe where. This is the most common position among organisations of your size, and the most expensive one to be in when something goes wrong, because the first time you map the exposure will be under pressure.

12–20 Drifting.

Adoption is running ahead of judgement. Tools have arrived faster than the decisions about them, and the governance you have is mostly held in the heads of two or three conscientious people. It works until one of them is on leave.

21–29 Deliberate.

You are making choices rather than absorbing them. The risk at this level is subtler: strong governance can coexist with automating precisely the capability that makes you worth hiring. Competence here does not protect you from the wrong sequence.

30–36 Governed.

Genuinely rare — and worth treating with suspicion. A high self-score is the most reliable predictor of a costly surprise, because judgement is the one faculty that cannot mark its own work. If you scored in this band, the questions to ask next are not on this page.

THE LIMIT OF A SELF-SCORE

What this cannot tell you.

Being honest about the instrument is part of using it well.

A diagnostic you score yourself is a mirror, not an examination. It is useful, and it is limited in three specific ways.

It cannot tell you what to do first.

A score of fourteen tells you that something is wrong. It does not tell you which of the four groups to address first, and the intuitive answer — start with the lowest subtotal — is usually the wrong one. Governance built before the judgement questions are settled tends to codify decisions nobody has actually made.

It cannot test whether your answers are true.

Every question here is answered from where you sit. In practice, the answers change considerably when the same questions are put to the people two levels below you, and the gap between those two sets of answers is frequently the most useful finding available about an organisation.

It cannot find the exposure you did not think to look for.

Twelve questions are twelve questions. They are the right ones to begin with, and they are not the whole of the ground. The material risks in an organisation are rarely where the low score sits — they are in the practice nobody thought to mention, because it has never been described as automation.

If this diagnostic has been uncomfortable in a useful way, that is the point of it, and it is as far as a document can honestly take you.

The practice.

The Human Premium is my advisory practice, working with founder-led businesses and social enterprises across the United Kingdom. I have run two companies as chief executive, co-founded a creative agency whose clients included Google, Meta, and The New York Times, and led organisations through acquisition and turnaround. The argument underneath the work is a simple one: as AI spreads, human judgement, culture, and trust become the things worth paying for.

WHAT COMES NEXT

The Assessment.

The same questions, asked beyond you.

The Assessment is the full version of what you have just done for yourself. It is a structured AI implementation readiness diagnostic conducted with you and with your organisation, and it does the three things this document cannot.

THIS DOCUMENT

Twelve questions, answered by one person.
A score, and a band.
An indication of where to look.
Free, and yours to keep.

THE ASSESSMENT

The same ground, tested against the people who do the work — not only the person who leads it.
A written finding, with the exposures named and evidenced.
A sequence: what to address first, what to defer, and what to stop.
A fixed fee, agreed before anything begins.

It is deliberately not a technology audit. Nobody will arrive to inventory your software. The work is to establish what your organisation should keep in human hands, what it can safely delegate, and in what order — so that adoption follows judgement rather than the other way round. It is the engagement most clients begin with, and it is designed to be useful whether or not anything follows it.

Begin a conversation.

Thirty minutes, no charge, and no sales process. Bring your score if you have one. We will work out whether The Assessment is the right next step, and I will tell you plainly if it is not.

info@thehumanpremium.co.uk

thehumanpremium.co.uk · Neural Pastoral, weekly on LinkedIn

Adopt what genuinely helps. Govern what genuinely risks. Keep what only humans do.

Built with AI. Thought by a human.