

BEYOND CHANGE READINESS

Data Readiness Is the New Change Readiness

Why Enterprise Transformation Now Needs the Data Assessment It Never Had

Giovanni Leonardi

August 2025



Data Readiness Is the New Change Readiness

Transformation · No. 355
August 2025

"Human operators have always been the error-correction layer between imperfect data and business outcomes."

1. The Green Dashboard That Hides the Real Risk

The change readiness slide is, by this point in the programme, reliably green. Stakeholder sentiment tracked in a traffic-light matrix, training completion rates climbing toward target, communications plans ticked off week by week. The discipline behind that slide is mature, well-funded, and largely unquestioned. If people are not ready for a new way of working, we know how to find out and what to do about it.

What nobody on the programme board is asking is whether the data is ready.

This was always a gap. It was never a crisis — until we started building systems that act on data autonomously. The current wave of agentic AI has done something no previous technology shift quite managed: it has made data quality a programme-stopping risk that sits outside every readiness framework we use. We have spent two decades refining how to assess whether people can absorb change. We have spent almost no time asking the equivalent question of the data estate, and the consequences are now arriving faster than most programme teams can absorb them.

2. The Asymmetry We Built and Never Noticed

Change readiness emerged because organisations learned, painfully, that technology deployments fail when people are not prepared. The discipline earned its place in the business case: budgets for training, dedicated change leads, readiness gates before go-live. Nobody now launches a major programme without it.

Data, by contrast, was always treated as infrastructure — present, presumed adequate, somebody else's problem. A data migration workstream might appear in the plan, but a systematic assessment of whether the data estate can support the transformation's intended outcomes? That barely exists as a concept in most programme methodologies.

The asymmetry was tolerable when systems were operated by people. A procurement officer working with a flawed supplier master knows which entries to trust and which to work around. A finance analyst running month-end close understands that the

cost-centre hierarchy has not been updated since the last restructure and compensates mentally. Human operators have always been the error-correction layer between imperfect data and business outcomes.

Agentic systems have no such layer. An agent tasked with reconciling purchase orders against contracts will act on whatever the supplier master tells it. If three customer records exist for the same entity with different credit terms, the agent will not pause to ask which is canonical — it will pick one, or worse, use all three at different points in the process. The failure is silent, confident, and compounding.

This is why the pattern emerging from early agent deployments is so consistent: the failures trace overwhelmingly to data, not to model capability. The agent works. The data does not support the work the agent is trying to do. And nobody assessed that gap before the programme committed its budget.

3. What the Parallel Teaches — and Where It Breaks

The change readiness discipline offers a surprisingly good structural template for what data readiness needs to become. Not because the domains are identical — they are not — but because the organisational mechanics are the same: an assessment that runs early, a set of criteria that gate progression, accountable ownership, and funded remediation.

The parallel breaks down in one important respect. Change readiness assesses a population's capacity to absorb a defined change. The population is known, the change is specified, and the gap can be measured through surveys, interviews, and observation. Data readiness must assess something more diffuse: whether a data estate — often sprawling, undocumented, and governed by convention rather than policy — can support outcomes that may themselves be emergent. An agent's interaction with data is not a single defined change; it is an ongoing, adaptive relationship. The assessment must therefore be more structural and less episodic than its change readiness equivalent.

With that distinction clear, the practical shape of a data readiness assessment comes into focus.

4. The Anatomy of a Data Readiness Assessment

A data readiness assessment answers five questions, each of which must be resolved before an agentic capability goes live:

1. Who owns this data — actually, not nominally? Most organisations have data ownership policies. Few have data ownership that functions. The assessment must identify, for every data domain the agent will touch, a named individual with genuine authority to define quality standards, approve access, and fund remediation. If that person cannot be found — and in many organisations, they cannot — the programme has its first finding before it has examined a single record.
2. What quality thresholds must the data meet, and does it meet them? This is not a generic data quality exercise. It is a targeted assessment against the specific

operations the agent will perform. If the agent will match invoices to contracts, the relevant question is whether contract records are complete, current, and unambiguous enough to support automated matching at the error rate the business will tolerate. The thresholds are defined by the use case, not by an abstract standard.

3. Is the lineage understood? Agents consuming data from multiple systems need lineage – the documented chain from source to the point of consumption. Without it, when an agent produces an unexpected output, there is no way to diagnose whether the fault lies in the model, the orchestration, or a data value that was transformed three times between its origin and the agent's input. In an agentic context, lineage is not a compliance nicety; it is an operational prerequisite for debugging.
4. Are access rules current and granular enough? Permissions in most enterprises reflect organisational structures that no longer exist. An agent operating under a service account inherits whatever access that account has been granted, which in practice often means either too much or too little. The assessment must verify that access controls are current, that they reflect the principle of least privilege, and that they have been designed for machine consumers rather than human ones.
5. What is the remediation path and what does it cost? Every gap the assessment surfaces must come with a costed remediation plan and a timeline. Data remediation is neither free nor fast. A programme that discovers in month eight that its customer master needs deduplication before agents can use it reliably has lost both time and credibility. The point of the assessment is to surface these costs early enough to fund them – or to descope the capability.

5. Who Owns It and When It Gates

The most common objection to treating data readiness as a gating criterion is that nobody wants to own it. Change readiness has a natural home – a change lead or change management office with a direct line to the programme director. Data readiness falls between the CTO's infrastructure teams, the CDO's governance function where one exists, and the programme's own delivery leadership.

The pattern I have watched work is straightforward: make data readiness the CDO's assessment and the programme director's gate. The CDO's team conducts the assessment because they hold the technical capability and the cross-domain view. The programme director owns the gate because they own the business case and the go/no-go decision. This mirrors exactly how change readiness operates: the change team assesses, the programme board decides.

Data readiness is the CDO's assessment and the programme director's gate – the same separation of assessment from decision that made change readiness work.

The gate sits at two points. The first is at business case approval: no agentic capability enters the portfolio without a preliminary data readiness assessment that identifies the data domains in scope, the current state of ownership and quality, and an order-of-magnitude remediation estimate. This is the point at which a programme

discovers that its ambition will cost twice what was assumed — or that the existing data estate makes it feasible.

The second gate is pre-deployment: a detailed assessment against the five questions above, with documented thresholds and sign-off from the data owner for every domain the agent will touch. This is the equivalent of the change readiness go/no-go that mature programmes run before go-live — the point at which the programme confirms that the data estate is fit for the autonomous operations it is about to enable.

6. The Business Case Implication

We would not approve a transformation programme that had no budget for change management. The current generation of agentic initiatives routinely approves programmes with no budget for data readiness — and the remediation costs surface later, as rework, delayed go-live, and agent failures that erode stakeholder confidence in the technology itself.

The economics of early discovery are not subtle. A customer master deduplication conducted during the design phase — when the data readiness assessment first flags the problem — is a contained exercise: a small team, four to six weeks, a defined scope. The same deduplication discovered in month eight, after the agent has been built against the assumption that customer records are unique, cascades. The agent logic needs reworking, the integration layer needs revisiting, the test cycles reset, and the programme board watches a go-live date that was already committed to the business slide quietly to the right. What was a funded remediation line becomes a programme crisis. The assessment exists to make the first version of that story happen instead of the second.

The discipline we need is not new. We built it for people readiness over twenty years, and the mechanics — assessment, ownership, gates, funded remediation — are well understood. What is new is the recognition that the agentic era has made data readiness a first-order programme risk, not an infrastructure assumption. The sooner that recognition reaches the business case template, the fewer programmes will discover their data problem after it has already become their delivery problem.