Assessment guide

# One process. Evaluated before anything is built, not after the platform is purchased.

A practical guide for evaluating AI opportunities before buying tools, building agents, or funding broad transformation programs.

> Axeron's assessment connects directly into implementation — the team that evaluates the opportunity is the team that builds and operates the first agent workflow, without a partner handoff.

ASSET PROMISE

**elps leaders move from vague AI ambition to a prioritized, measurable first use c ase.**

BEST AUDIENCE

CEOs, COOs, CIOs, transformation leaders, government modernization teams

RECOMMENDED USE

Use before or after an enterprise briefing to help the buyer align internal stakeholders and define a practical next step.

METHODOLOGY

The AI Transformation Assessment follows a structured six-dimension sequence. Each dimension is evaluated before a use case is recommended. Skipping a dimension means the pilot that follows is built on assumption rather than evidence about the operating environment.

01

### Map how work actually happens

The assessment begins with the process, not the platform. Axeron reviews how work flows through the organization — including handoffs between teams, manual steps, approval loops, rework cycles, and informal decisions that exist nowhere in documentation. The goal is to understand the real operating model before recommending where AI should enter it.

02

### Measure the current baseline

Before any AI recommendation, the assessment establishes observable facts about current performance: volume per period, average cycle time, known error or rework rate, manual effort per case. Qualitative estimates from interviews are acceptable when system data is unavailable. The baseline is what makes a pilot result credible to a CFO.

03 **Review data quality and source-of-truth gaps**

Axeron maps what data the candidate process runs on — which systems hold it, how clean it is, who maintains it, and whether an agent can reliably access it. Data gaps are surfaced as pre-conditions that must be resolved before a pilot can produce valid results, not discovered during build.

04 **Evaluate governance and deployment constraints**

Every organization has constraints that shape which AI architecture is viable: data residency requirements, regulated decision points, human approval obligations, security boundaries, and deployment preferences ranging from SaaS to fully dark environments. These are evaluated explicitly — not assumed — so the recommended approach fits within the actual operating rules.

05 **Score candidate use cases**

Candidate processes are scored on two axes: business value and AI readiness. The scoring is ordinal and qualitative — it ranks use cases against each other rather than generating a precise numeric score. The output is a prioritized map that separates viable first pilots from ideas that are high value but not yet ready, technically feasible but operationally low-impact, or both weak and unready.

06 **Define the pilot recommendation**

The assessment closes with one concrete recommendation: a named first use case, a defined success measure grounded in the observed baseline, a summary of the governance model the pilot requires, and a handoff summary that feeds directly into Process Discovery and AxeStudio implementation. The output is designed to support a budget conversation, not replace one.

## 01 Why this guide exists

Most organizations do not fail because they lack AI ideas. They fail because they start with a tool before they understand the process, the data, the people, the controls, and the economic case. Axeron's assessment starts with how work actually happens, then determines where AI can create measurable operational value.

## 02 What Axeron evaluates

The assessment looks at the full operating environment, not just model readiness.

- Business processes and decision points
- Manual effort and rework loops
- Documents, forms, emails, and system handoffs
- Data quality and source-of-truth gaps
- Security, sovereignty, and deployment constraints
- Stakeholder readiness and change-management risk
- Governance, auditability, and human approval requirements

## 03 Assessment output

The buyer receives a ranked AI opportunity map that separates high-value first use cases from ideas that are too risky, too vague, or not ready. The output is designed for internal leadership alignment and budget discussion.

- Current-state workflow map
- Future-state AI-enabled workflow
- Use-case priority matrix
- Data and integration readiness view
- Governance and risk notes
- Pilot recommendation
- KPIs and value hypothesis

## 04 ow Axeron is different

Axeron does not only advise. The assessment connects directly into Process Discovery, AxeStudio implementation, TeamAI workflows, sovereign deployment, MLOps/LLMOps, and Continuum governance. The strategy can become a working system without changing partners halfway through.

## 05 Best first use case pattern

The strongest first use cases are not the flashiest. They usually involve repeated document work, slow reviews, operational handoffs, high-volume requests, or decisions that require evidence and approval.

- Clear business owner
- Measurable baseline
- Repeatable workflow
- Available data
- Manageable risk boundary
- Visible cost, time, quality, or compliance impact

AI-OPPORTUNITY MATRIX

The Use-Case Priority Matrix plots candidate AI opportunities on two qualitative axes — business value and AI readiness — so leadership can align on where to start without relying on invented benchmarks. Placement is directional and based on the six evaluation dimensions. The purpose is to force a structured conversation about trade-offs, not to replace one.

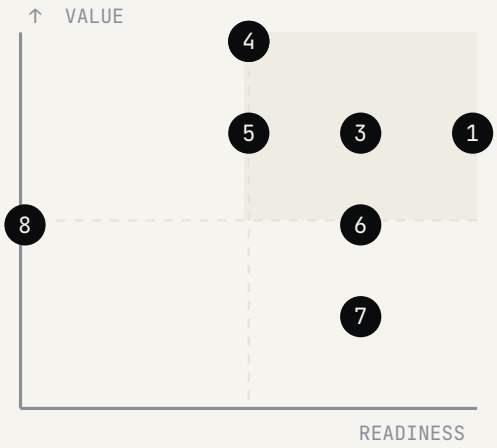


- Document intake and extraction
- Compliance and audit evidence preparation
- High-volume operational approvals
- Due diligence and review preparation
- Citizen service triage and routing
- Procurement and contract review
- Ad hoc executive reporting requests
- Bespoke strategic analysis

### Start here

High value, high readiness. The process has a clear owner, a measurable baseline, structured input documents, available data, and a containable risk boundary. This is the first pilot.

### Prepare first

High value, lower readiness. The business case is strong but data quality, ownership clarity, or governance design must be resolved before a pilot can succeed. The assessment defines what to fix.

### Consider later

High readiness, lower value. The process is technically tractable but the operational impact is modest relative to the cost of change. Useful as a confidence-building exercise after the first pilot.

### Deprioritize

Low value, low readiness. Irregular, contested, or data-poor. Returning to this after the organization has built trust in the first working system is the right sequence.

## WORKFLOW-READINESS SCORECARD

Run this rubric against one candidate use case before recommending a pilot. Score each dimension as Weak, Workable, or Strong. A mix of Workable and Strong across all six means the process is a viable first pilot. A single Weak on Business Ownership or Risk Boundary is a blocker — resolve it before any build work begins.

01 **Business ownership**

Is there one named person accountable for this process end to end — someone who can approve the future-state design, own the pilot KPIs, and handle exceptions when the system is wrong? Strong: named owner with budget authority and operational decision rights. Workable: named owner with shared accountability in one adjacent team. Weak: ownership is a committee, informally distributed, or actively contested.

02 **Baseline measurability**

Can the current performance of this process be observed, even approximately? Strong: volume, cycle time, and error rate are already tracked in a system. Workable: these can be estimated from interviews and system logs within a week. Weak: no one can state how much work flows through, how long it takes, or how often it goes wrong.

03 **Repeatability and pattern stability**

Does this process follow a recognizable pattern most of the time, or is every instance bespoke? Strong: the same steps apply to the same inputs in the same order for more than 80% of cases; exceptions are categorizable. Workable: a core pattern exists with a manageable set of exception types. Weak: every case is negotiated from scratch or depends heavily on tacit judgment with no shared rules.

04 **Data clarity and access**

Are the inputs to this process available, reasonably clean, and traceable to a known system? Strong: source documents and data are in defined locations; a person can find them consistently without manual assembly. Workable: sources are known but scattered across two or three systems and require manual consolidation. Weak: data is missing, siloed, of unknown quality, or held in systems with no accessible export path.

05 **Risk boundary definition**

Can the high-stakes decision or action in this process be cleanly separated from the preparatory work an agent would handle? Strong: there is a clear, policy-backed approval step a human owns; the agent prepares but does not decide. Workable: a boundary can be designed with moderate governance and change-management effort. Weak: human judgment and automated action are entangled throughout, with no natural checkpoint a policy or audit requirement enforces.

06 **Visible operational impact**

If this process ran faster, more consistently, or with fewer errors, would leadership notice and would that improvement connect to a cost center, compliance obligation, or service-level commitment? Strong: the impact is direct and connects to a known problem a sponsor already owns. Workable: the impact is real but affects an internal team rather than an external stakeholder or regulatory obligation. Weak: the value of improvement is diffuse, contested, or depends on assumptions no one has tested.

## WHY AI PILOTS FAIL

The most common AI transformation failures are not technology failures. They are sequencing failures — the organization commits to a platform, a vendor, or a scope before it understands the process, the data, or the governance requirement. Each pattern below is recognizable before a pilot is designed.

### WHY AI PILOTS FAIL

- × Do not select a platform before completing the assessment. The platform selection should follow from the process map, the data environment, the governance constraints, and the deployment requirement — not precede them. Buying first anchors every subsequent decision to what the tool can do rather than what the work requires.
- × Do not choose the most visible or politically significant process as the first pilot. The first pilot should be chosen for tractability — bounded scope, structured documents, clear ownership, and a manageable risk boundary. A high-profile failure on a contested process sets back the entire program; a successful pilot on a bounded one builds the organizational trust the harder cases require.
- × Do not define pilot success without a pre-existing baseline. A metric that is invented after the pilot begins cannot be used to claim improvement. If cycle time, volume, or error rate is not observable before the pilot, estimate it from interviews and document it — then hold to that number regardless of what the pilot produces.
- × Do not proceed to build when business ownership is unclear. An AI workflow that no one is accountable for will drift, accumulate exceptions, and be quietly abandoned. Disputed or committee ownership is not a governance nuance — it is a hard blocker. Resolve it before the assessment recommends a pilot.
- × Do not treat governance and deployment constraints as implementation details. Data residency requirements, regulated decision points, human approval obligations, and security boundaries must be understood before a use case is recommended. Discovering a constraint during build is expensive; discovering it after launch is a program-ending event.

### WHAT TO DO INSTEAD

- → Start the assessment with one painful, observable process — not a transformation vision. A COO who can name one process, measure its current cost, and describe its failure mode is in a better position to commission a pilot than one who has a broad AI strategy without a specific starting point.

- → Confirm a named, accountable process owner before any discovery or design work begins. That person must have the authority to approve the future-state workflow, own the pilot KPIs, and escalate exceptions when the system behaves unexpectedly.
- → Establish an observable baseline — even a rough estimate from interviews — before defining the pilot success metric. The baseline is what makes the improvement credible to a finance team and defensible in a post-pilot review.
- → Evaluate governance and deployment constraints as first-class assessment inputs. Whether the organization requires data residency, on-premises deployment, human approval for every consequential action, or a full audit trail should be known before a use case is scoped — not discovered when an architecture decision needs to be reversed.
- → Choose the first use case for its tractability, then use its success to build organizational appetite for the more complex opportunities that scored high on value but low on readiness.

## WHAT YOU WALK AWAY WITH

At the close of the AI Transformation Assessment, the organization receives a structured set of working documents. Each output is designed for a specific internal audience: the CEO or COO defending the investment decision, the CIO evaluating architecture requirements, the CFO approving the pilot budget, and the implementation team that will design and build the first agent workflow.

01 Current-state workflow map for the candidate use case — showing actual process steps, system handoffs, approval loops, exception paths, and rework cycles, annotated with the interview and document sources that informed each element.

02 Use-case priority matrix plotting candidate processes on business value and AI readiness axes, with qualitative placement rationale for each — giving leadership a prioritization tool that does not depend on invented benchmarks or vendor-supplied figures.

03 Readiness assessment completed against six dimensions for the top one or two candidate processes, with a plain interpretation of what Workable and Weak ratings mean for the pilot scope and what must be resolved before build can begin.

04 Data and integration readiness notes identifying source systems, document locations, known data quality gaps, access constraints, and any integration dependencies the platform deployment will need to address.

05 Governance and deployment summary covering which steps require human approval, what audit fields the workflow must generate, what data residency or sovereignty constraints apply, and what deployment model the process's risk profile requires.

06 Pilot recommendation with a named first process, a defined success measure grounded in the observed baseline, a summary of the governance model the pilot requires, and a handoff document that feeds directly into Process Discovery and AxeStudio implementation.

## WHERE IT FEEDS

The AI Transformation Assessment is the entry point into the full Axeron operating model. The assessment output feeds directly into Process Discovery, which maps the selected workflow in detail and designs the AI-enabled future state. That design feeds into AxeStudio, which builds and implements the production workflow. The Axeron platform handles deployment, integration, and runtime governance through Continuum — monitoring agent behavior, gating self-improvement, and maintaining the audit trail. Data and model readiness work runs in parallel when the candidate process requires fine-tuning or structured document preparation before an agent can operate reliably. The design constraint throughout is continuity: the team that evaluated the opportunity is the team that maps the process, builds the agent, and governs the system in production. That constraint exists because the assessment findings — the exception

logic, the undocumented decision rules, the governance boundaries — must make it into the production system intact, not be reinterpreted by a separate implementation partner who was not present when they were discovered.

BEFORE YOU START

- [ ] Name one painful process
- [ ] List systems and documents involved
- [ ] Define what cannot be automated
- [ ] Measure current volume and cycle time
- [ ] Identify who approves the work
- [ ] Choose one measurable pilot outcome

NEXT STEP

## Start with an AI Transformation Assessment, not a generic platform demo.

Axeron turns resource-level education into a scoped conversation: one process, one measurable outcome, one deployment model, and one governance path.