

The AI Investment File

How Leaders Find, Buy, and Implement AI That Pays

Ugur Ozer

The Argument in Brief

A complete overview from the book. Original pagination is retained. This sample includes the publication notice and the selected overview only.

Books, purchase links and free working templates

airiskmanagement.ca/books

Publication Notice

Copyright 2026 Ugur Ozer. All rights reserved.

Published by Aeon Digital Publishing.

Sources updated through 2026.

This book provides general management information, not legal, investment, security, or other professional advice. Laws, standards, products, and technical capabilities change. Verify current requirements and evidence for the system and jurisdiction at issue.

Public cases are cited. Scenarios are labeled by type. Illustrative figures are not observed client results.

All Leadership Principles are original author text. Historical quotations are used for perspective, not as evidence.

The Argument in Brief

The Short Version

Whether your company employs ten people or ten thousand, it probably has more AI activity than you can see and less AI value than the activity suggests.

An employee uses an assistant to rewrite an email while a marketing contractor creates images. A salesperson summarizes a call, and a software feature quietly classifies a customer request. Elsewhere, someone uploads a sensitive document to a free tool. A vendor describes an ordinary automation feature as "AI-powered." All of these may count as AI use. None proves that AI has improved the business.

The leader's job is not to maximize the number of AI tools, pilots, or use cases. It is to make a small number of good investment decisions and to demand evidence before dependence.

The SME AI decision funnel

Start with business friction. Earn the right to discuss a tool.




Figure 3. The SME AI decision funnel: many observations of work become a short list of opportunities, one controlled test, and only then a scale decision.

What this figure shows: The funnel starts wide, with the many observations of work a leader can collect in the first 30 days, and narrows step by step. For the service-classification teaching Receipt in this section, the everyday queue of customer requests becomes one ranked task (classify by service type), then one isolated capability test with employee confirmation, and only then a scale decision leadership can make on evidence. Most of what enters the top never reaches the bottom, and that is the point: the funnel exists to protect attention, not to feed a tool count.

This book gives you one operating framework for doing that: **RECEIPTS**.

- **R - Result:** What business result should change, for whom, and by how much?
- **E1 - Economics:** What will the system really cost, how will value be captured, and when might it pay back?
- **C - Capability and control:** Can it perform the actual task, where does it fail, and what limits its action?
- **E2 - Evidence:** What baseline, comparison, and source support the claim?
- **I - Information:** What data enters, leaves, persists, trains, or crosses a boundary?
- **P - People, process, and permissions:** Who owns, checks, overrides, and remains accountable?
- **T - Tracking:** What will be monitored after release, including quality, cost, reliance, drift, and incidents?
- **S - Stop or scale:** What threshold triggers expansion, redesign, rollback, replacement, or retirement?

You will build one **AI Investment Receipt** for every material candidate. The **AI Investment File** is the organization's Book 1 container. Each Receipt is a single decision record for one candidate, and each chapter adds a supporting artifact to it. Chapter 12 assembles the Receipts into the **90-Day AI Portfolio**.

If a proposal cannot answer the RECEIPTS questions, it is not ready for approval. It may still deserve a small experiment, but it does not deserve belief.

Five Ideas to Keep

1. Find the work before you find the tool

Do not begin with "Where can we use generative AI?" Begin with "Where does valuable work get stuck?" Look for queues, rework, repeated search, and inconsistent decisions. Also look for abandoned

demand, expert bottlenecks, slow handoffs, and avoidable errors. Then ask whether AI is one plausible way to improve that workflow.

2. Test the real task, not the demo

AI capability is uneven. A system can look brilliant on a polished example and fail on the awkward cases that dominate your operating cost. Test it on representative work from your business, including ordinary cases, difficult cases, rare cases, and cases where a plausible mistake would be costly.

3. Time saved is not automatically money earned

If a tool saves time, what operating constraint changes? The team might shorten a queue, serve more customers, or improve quality. Nothing may change at all. The economic result comes from the changed workflow, not from the stopwatch alone.

4. Buying is an architecture decision

The cheapest subscription may create expensive integration, review, security, or switching costs. The apparently difficult private or customized option may be justified when proprietary knowledge, customer information, reliability, or strategic differentiation matters. Decide what the company can rent, what it must configure, and what it needs to control.

5. Set the stop rule before enthusiasm takes over

Every pilot needs a baseline, a success threshold, a budget boundary, and a date for a decision. A stopped pilot is not necessarily a failure. It can be a cheap and useful answer. One expensive outcome is an indecisive pilot: it produces too little evidence to scale but lingers because nobody ends it.

Your First 30 Days

Do not launch a transformation program this month. Establish control of the facts.

1. List every AI tool and AI-enabled feature that teams knowingly use.
2. Identify the workflow and task supported by each one.
3. Record who uses it, who approved it, and who owns the result.
4. Mark the information it receives: public, internal, confidential, personal, regulated, or strategically sensitive.
5. Classify its operating level: informal experiment, authorized assistance, embedded workflow, or material reliance.
6. Record what evidence exists: anecdote, self-reported time saving, measured task improvement, captured business value, or no evidence.
7. Stop duplicate, abandoned, or obviously inappropriate uses. Do not pretend the remaining list is a strategy.
8. Choose one workflow for a deeper opportunity review in Chapter 2.

At the end of the month, you should be able to answer one question without calling a vendor. **Where are we already relying on AI, and what do we actually know about the result?**

The One-Sentence Rule

For any proposed AI investment, require the sponsor to complete this sentence:

We will use [system] to help [people] perform [specific task] in [workflow], so that [business measure] changes from [baseline] to [target] within [time], while [named owner] remains accountable and we stop or redesign if [threshold] is breached.

If the sentence cannot be completed, the proposal is not yet an investment case. It is an idea.

A Completed Illustrative Receipt

Scenario.

It previews the continuing service-classification scenario used later in the book. Unknowns remain visible instead of being completed with assumptions.

- **Receipt identifier:** Service-classification teaching Receipt
- **Version:** Current illustrative state
- **Current gate:** Capability test complete; business case not begun
- **Decision date:** Not yet assigned

RECEIPTS field	Completed illustrative entry
R - Result	In customer-service intake, suggest a service type for eligible ordinary requests and require employee confirmation. The intended result is fewer wrong routes and faster first useful action. A local baseline exists in the scenario; a target and causal effect are not yet evidenced.
E1 - Economics	No return is claimed. Full cost, usable capacity, capture mechanism, learning ceiling, and payback threshold are unknown or not yet evidenced.
C - Capability and control	The illustrative test shows a useful zone for ordinary cases, difficulty with multi-issue requests, and a missed safety cue. Employee confirmation remains mandatory. Automated priority, replies, scheduling, and high-consequence cases are outside scope.
E2 - Evidence	The findings are illustrative capability evidence within this scenario only. They are not organizational results. Live workflow effect, attribution, and captured value are not yet evidenced.

RECEIPTS field	Completed illustrative entry
I - Information	Customer requests are the proposed input. Prohibited attachments must stay outside the external route. Supplier processing, retention, downstream services, and the final business record are not yet evidenced.
P - People, process, and permissions	The service leader owns routing and escalation. An employee confirms every suggestion and routes excluded cases. Peak-period reviewer capacity is unknown and must be tested.
T - Tracking	Candidate measures are wrong routing, time to first useful action, missed escalation, recontact, review effort, information events, and queue recovery. Thresholds and response owners are not yet set.
S - Stop or scale	Advance only to a bounded business-case and readiness review. Do not begin live use. Stop or redesign if the escalation, information, record, fallback, or review-capacity conditions cannot be supported. The next decision date remains pending.

This Receipt is complete for its current decision because every field is evidenced, forecast, explicitly unknown, or pending. Complete does not mean favorable, proven, or ready for live use.

Decide, Operate, Prove

Book 1 creates the investment record. Later volumes add operating and assurance evidence without replacing the decision history.

SERIES MAP

One leadership journey, three decisions

Each book adds evidence to the same control file.




Figure 1. The series journey: decide what deserves investment, operate the resulting human and digital workforce, then prove that material systems remain trustworthy.

ONE SHARED DECISION RECORD

Eight Questions, One Receipt

RECEIPTS keeps the business result, full cost, evidence, information, authority, monitoring, and decision in one inspectable record.

BOOK 1 · DECIDE

RECEIPTS: eight questions, one decision

A credible use case answers every field.




Figure 2. RECEIPTS turns one AI proposal into eight linked decisions. Weakness in any field changes the investment case.