

The AI Buyer's Field Guide

Choose, Contract, Test, and Govern AI Vendors

Ngo Trung Kien

Contents

Copyright and publication notice	1
Introduction: the promise and the boundary	3
Visual frontispiece	4
0.1 Who this book is for	4
0.2 Thesis and falsification condition	4
1 Buy an Outcome, Not AI	7
1.1 In plain terms	7
1.1.1 Chapter contract	7
1.2 The decision at stake	7
1.3 Working model	8
1.3.1 See the chapter in practice	8
1.3.2 Chapter-specific operating protocol	10
1.3.3 Operating workflow	12
1.3.4 Evidence and control graph	13
1.4 Worked composite case file — illustrative, not empirical evidence	14
1.4.1 Scenario trace	15
1.4.2 Artifact handoff	16
1.4.3 Evidence application	16
1.4.4 Transfer test	16
1.5 Measurement and recovery decision	16
1.6 Practice lab	17
1.7 Limits and judgment	17
1.8 Takeaways	18
1.9 Reflection	18
2 Classify Consequence and Exposure	19
2.1 In plain terms	19
2.1.1 Chapter contract	19
2.2 The decision at stake	19
2.3 Working model	20
2.3.1 See the chapter in practice	20
2.3.2 Chapter-specific operating protocol	22
2.3.3 Operating workflow	24
2.3.4 Evidence and control graph	25
2.4 Worked composite case file — illustrative, not empirical evidence	26
2.4.1 Scenario trace	27
2.4.2 Artifact handoff	28
2.4.3 Evidence application	28
2.4.4 Transfer test	28

2.5	Measurement and recovery decision	29
2.6	Practice lab	29
2.7	Limits and judgment	30
2.8	Takeaways	30
2.9	Reflection	30
3	Write Disqualifying Requirements	31
3.1	In plain terms	31
3.1.1	Chapter contract	31
3.2	The decision at stake	31
3.3	Working model	32
3.3.1	See the chapter in practice	32
3.3.2	Chapter-specific operating protocol	34
3.3.3	Operating workflow	36
3.3.4	Evidence and control graph	37
3.4	Worked composite case file — illustrative, not empirical evidence	38
3.4.1	Scenario trace	39
3.4.2	Artifact handoff	40
3.4.3	Evidence application	40
3.4.4	Transfer test	40
3.5	Measurement and recovery decision	41
3.6	Practice lab	41
3.7	Limits and judgment	42
3.8	Takeaways	42
3.9	Reflection	42
4	Ask for Evidence, Not Adjectives	43
4.1	In plain terms	43
4.1.1	Chapter contract	43
4.2	The decision at stake	43
4.3	Working model	44
4.3.1	See the chapter in practice	44
4.3.2	Chapter-specific operating protocol	46
4.3.3	Operating workflow	48
4.3.4	Evidence and control graph	49
4.4	Worked composite case file — illustrative, not empirical evidence	50
4.4.1	Scenario trace	51
4.4.2	Artifact handoff	52
4.4.3	Evidence application	52
4.4.4	Transfer test	52
4.5	Measurement and recovery decision	53
4.6	Practice lab	53
4.7	Limits and judgment	54
4.8	Takeaways	54
4.9	Reflection	54

5	Map the Supply Chain	55
5.1	In plain terms	55
5.1.1	Chapter contract	55
5.2	The decision at stake	55
5.3	Working model	56
5.3.1	See the chapter in practice	56
5.3.2	Chapter-specific operating protocol	58
5.3.3	Operating workflow	60
5.3.4	Evidence and control graph	61
5.4	Worked composite case file — illustrative, not empirical evidence	62
5.4.1	Scenario trace	63
5.4.2	Artifact handoff	64
5.4.3	Evidence application	64
5.4.4	Transfer test	64
5.5	Measurement and recovery decision	65
5.6	Practice lab	65
5.7	Limits and judgment	66
5.8	Takeaways	66
5.9	Reflection	66
6	Design a Comparable Evaluation	67
6.1	In plain terms	67
6.1.1	Chapter contract	67
6.2	The decision at stake	67
6.3	Working model	68
6.3.1	See the chapter in practice	68
6.3.2	Chapter-specific operating protocol	70
6.3.3	Operating workflow	72
6.3.4	Evidence and control graph	73
6.4	Worked composite case file — illustrative, not empirical evidence	74
6.4.1	Scenario trace	75
6.4.2	Artifact handoff	76
6.4.3	Evidence application	76
6.4.4	Transfer test	76
6.5	Measurement and recovery decision	77
6.6	Practice lab	77
6.7	Limits and judgment	78
6.8	Takeaways	78
6.9	Reflection	78
	Midpoint visual synthesis	79
7	Plan Human Oversight and Appeals	81
7.1	In plain terms	81
7.1.1	Chapter contract	81
7.2	The decision at stake	81
7.3	Working model	82

7.3.1	See the chapter in practice	82
7.3.2	Chapter-specific operating protocol	84
7.3.3	Operating workflow	86
7.3.4	Evidence and control graph	87
7.4	Worked composite case file — illustrative, not empirical evidence	88
7.4.1	Scenario trace	89
7.4.2	Artifact handoff	90
7.4.3	Evidence application	90
7.4.4	Transfer test	90
7.5	Measurement and recovery decision	90
7.6	Practice lab	91
7.7	Limits and judgment	91
7.8	Takeaways	92
7.9	Reflection	92
8	Contract for the Lifecycle	93
8.1	In plain terms	93
8.1.1	Chapter contract	93
8.2	The decision at stake	93
8.3	Working model	94
8.3.1	See the chapter in practice	94
8.4	Contract issue prompts	96
8.4.1	Chapter-specific operating protocol	96
8.4.2	Operating workflow	99
8.4.3	Evidence and control graph	100
8.5	Worked composite case file — illustrative, not empirical evidence	101
8.5.1	Scenario trace	102
8.5.2	Artifact handoff	103
8.5.3	Evidence application	103
8.5.4	Transfer test	103
8.6	Measurement and recovery decision	103
8.7	Practice lab	104
8.8	Limits and judgment	104
8.9	Takeaways	105
8.10	Reflection	105
9	Pilot Under Real Boundaries	107
9.1	In plain terms	107
9.1.1	Chapter contract	107
9.2	The decision at stake	107
9.3	Working model	108
9.3.1	See the chapter in practice	108
9.3.2	Chapter-specific operating protocol	110
9.3.3	Operating workflow	112
9.3.4	Evidence and control graph	113
9.4	Worked composite case file — illustrative, not empirical evidence	114

9.4.1	Scenario trace	115
9.4.2	Artifact handoff	116
9.4.3	Evidence application	116
9.4.4	Transfer test	116
9.5	Measurement and recovery decision	116
9.6	Practice lab	117
9.7	Limits and judgment	117
9.8	Takeaways	118
9.9	Reflection	118
10	Monitor What Can Change	119
10.1	In plain terms	119
10.1.1	Chapter contract	119
10.2	The decision at stake	119
10.3	Working model	120
10.3.1	See the chapter in practice	120
10.3.2	Chapter-specific operating protocol	122
10.3.3	Operating workflow	124
10.3.4	Evidence and control graph	125
10.4	Worked composite case file — illustrative, not empirical evidence	126
10.4.1	Scenario trace	127
10.4.2	Artifact handoff	128
10.4.3	Evidence application	128
10.4.4	Transfer test	128
10.5	Measurement and recovery decision	128
10.6	Practice lab	129
10.7	Limits and judgment	129
10.8	Takeaways	130
10.9	Reflection	130
11	Prepare Incident and Exit Paths	131
11.1	In plain terms	131
11.1.1	Chapter contract	131
11.2	The decision at stake	131
11.3	Working model	132
11.3.1	See the chapter in practice	132
11.3.2	Chapter-specific operating protocol	134
11.3.3	Operating workflow	136
11.3.4	Evidence and control graph	137
11.4	Worked composite case file — illustrative, not empirical evidence	138
11.4.1	Scenario trace	139
11.4.2	Artifact handoff	140
11.4.3	Evidence application	140
11.4.4	Transfer test	140
11.5	Measurement and recovery decision	140
11.6	Practice lab	141

11.7 Limits and judgment 141

11.8 Takeaways 142

11.9 Reflection 142

12 Make the Renewal Decision Anew 143

12.1 In plain terms 143

12.1.1 Chapter contract 143

12.2 The decision at stake 143

12.3 Working model 144

12.3.1 See the chapter in practice 144

12.3.2 Chapter-specific operating protocol 146

12.3.3 Operating workflow 148

12.3.4 Evidence and control graph 149

12.4 Worked composite case file — illustrative, not empirical evidence150

12.4.1 Scenario trace 151

12.4.2 Artifact handoff 152

12.4.3 Evidence application 152

12.4.4 Transfer test 152

12.5 Measurement and recovery decision 152

12.6 Practice lab 153

12.7 Limits and judgment 153

12.8 Takeaways 154

12.9 Reflection 154

Closing: keep the evidence visible 155

A compact implementation workbook 157

Decision register 158

Evidence scorecard 159

Exception log 160

Ninety-day plan 161

Control and test matrix 162

Workflow map 163

13 References 165

Buy an Outcome, Not AI

In plain terms

A procurement starts with a measurable use case, affected people, and failure cost rather than a technology label. In an AI procurement and contract decision, the practical question is **What business condition should change if the purchase succeeds?** For business, procurement, legal, security, risk, and product leaders buying AI-enabled services, the answer counts only when **Write an outcome statement with baseline, target, users, decisions, and non-goals** produces reviewable evidence for **outcome baseline established** without normalizing **issuing a request for a generic AI platform**.

Chapter contract

DECIDE

What business condition should change if the purchase succeeds?

BUILD

Write an outcome statement with baseline, target, users, decisions, and non-goals.

MEASURE

outcome baseline established

STOP OR REDESIGN WHEN

issuing a request for a generic AI platform

The decision at stake

For **Buy an Outcome, Not AI**, the procurement owner with legal and security review must decide **what business condition should change if the purchase succeeds?** before the buying decision crosses from discussion into consequence for buyers, users, affected data subjects, and vendors. The minimum decision record is the artifact **Write an outcome statement with baseline, target, users, decisions, and non-goals**, the observable **outcome baseline established**, and the explicit hold condition **issuing a request for a generic AI platform**. If any Chapter 1 field is missing, the safe default is to make no commitment and preserve the unresolved requirement.

The bounded **Buy an Outcome, Not AI** decision advances the book's promise to turn AI procurement from a feature comparison into a lifecycle

decision backed by evidence, contract terms, tests, and exit options. It does not endorse a product or predetermine the Chapter 1 answer; it creates a record that another person can challenge at each evaluation gate, contract change, and renewal review.

Working model

Read the **Buy an Outcome, Not AI** model as four linked fields. **Outcome:** A procurement starts with a measurable use case, affected people, and failure cost rather than a technology label. **Boundary:** the design must make **issuing a request for a generic AI platform** visible rather than convenient. **Evidence:** interpret **outcome baseline established** together with review effort and excluded cases. **Authority:** the procurement owner with legal and security review may continue, narrow, escalate, or stop the buying decision. For **Buy an Outcome, Not AI**, Reference [1] from UK Government, *Guidelines for AI procurement*, informs a Chapter 1 boundary; it is not evidence that the local method or illustrative target will work.

See the chapter in practice

The **buy an outcome, not ai in practice** maps *The AI Buyer's Field Guide* Chapter 1 into a field scene before its abstract models begin. In one frame, it holds the question **What business condition should change if the purchase succeeds?**, the artifact **Write an outcome statement with baseline, target, users, decisions, and non-goals.**, the proof signal **outcome baseline established**, and the warning **issuing a request for a generic AI platform**.




Figure 1.1: Chapter 1 field scene for Buy an Outcome, Not AI: a boundary view in the book's buyer setting.

Read the scene: For The AI Buyer's Field Guide, Chapter 1, begin with the central buyer scene and scan DECIDE, MAKE, CHECK, and WATCH. Treat the Buy an Outcome, Not AI scene as conceptual, then use its cases and evidence records to test whether outcome baseline established actually improves. For **Buy an Outcome, Not AI**, this explanatory plate is a reasoning prompt—not observed evidence or a universal operating prescription.

Visual field note — Buy an Outcome, Not AI:

- 1. Locate the tension.** In this composite situation, A public-service team reframes a chatbot purchase as reducing unresolved routine inquiries without blocking human access.
- 2. Name the choice.** Answer the chapter question: **What business condition should change if the purchase succeeds?**
- 3. Produce the record.** Use **Write an outcome statement with baseline, target, users, decisions, and non-goals.**, then retain the procurement owner, assumptions, and case evidence needed to review it.
- 4. Test and return.** Track **outcome baseline established**. If **issuing a request for a generic AI platform**, return to DECIDE and revise the contract boundary instead of explaining the warning away.

Chapter-specific operating protocol

Phase	Reader action	Evidence retained	Hold condition
1. Define	Write the exact decision behind What business condition should change if the purchase succeeds?	Chapter 1 owner, input, output, exception, and timestamp.	Hold Buy an Outcome, Not AI if its record cannot be reconstructed.
2. Contrast	Apply Write an outcome statement with baseline, target, users, decisions, and non-goals to one ordinary case.	Chapter 1 owner, input, output, exception, and timestamp.	Hold Buy an Outcome, Not AI if its record cannot be reconstructed.
3. Exclude	Run the same artifact against a counter case shaped by issuing a request for a generic AI platform.	Chapter 1 owner, input, output, exception, and timestamp.	Hold Buy an Outcome, Not AI if its record cannot be reconstructed.
4. Assign	Name the procurement owner with legal and security review as the person who may continue, narrow, or stop.	Chapter 1 owner, input, output, exception, and timestamp.	Hold Buy an Outcome, Not AI if its record cannot be reconstructed.
5. Challenge	Record outcome baseline established and the review effort for the same case.	Chapter 1 owner, input, output, exception, and timestamp.	Hold Buy an Outcome, Not AI if its record cannot be reconstructed.

Phase	Reader action	Evidence retained	Hold condition
6. Publish	Release the Buy an Outcome, Not AI artifact only when its evidence answers this chapter’s question.	Chapter 1 owner, input, output, exception, and timestamp.	Hold Buy an Outcome, Not AI if its record cannot be reconstructed.

This sequence is specific to **Buy an Outcome, Not AI**. For that Chapter 1 decision, it is not permission to automate; it is a way for business, procurement, legal, security, risk, and product leaders buying AI-enabled services to make the reasoning inspectable at each evaluation gate, contract change, and renewal review.

Write an outcome statement with baseline, target, users, decisions, and non-goals. This is the chapter’s core production artifact. The **Buy an Outcome, Not AI** artifact should fit on one page, use ordinary language, and be reviewed by someone who did not design it. A compact artifact exposes disagreements that a long narrative can conceal.

The chapter’s visual tool is the **chapter contract**. For **Buy an Outcome, Not AI**, it translates explanation into an inspectable decision aid: the reader can see the decision, artifact, evidence, and exact condition that pauses the buying decision.




Figure 1.2: Chapter contract connecting the decision, artifact, measure, and stop condition for Buy an Outcome, Not AI.

How to use this page: Move across the four Buy an Outcome, Not AI cards; compare the decision with outcome baseline established and stop before issuing a request for a generic AI platform. The chapter-specific measure is **outcome baseline established**; the principal warning is **issuing a request for a generic AI platform**.

Operating workflow

The **buy an outcome, not ai workflow** makes the chapter executable. For **Buy an Outcome, Not AI**, the five stages carry **Write an outcome statement with baseline, target, users, decisions, and non-goals**. from the question **What business condition should change if the purchase succeeds?** to a test against **outcome baseline established**, with **issuing a request for a generic AI platform** as an explicit return or stop signal.




Figure 1.3: Five-stage workflow for chapter 1, Buy an Outcome, Not AI: frame the decision, set the contract boundary, build Write an outcome statement with baseline, target, users, decisions, and non-goals., test against outcome baseline established, then continue, narrow, or stop before issuing a request for a generic AI platform.

Read the workflow: Move from FRAME to DECIDE for Buy an Outcome, Not AI. The lower return path is mandatory when evidence is missing or the stop signal—issuing a request for a generic AI platform—appears. For **Buy an Outcome, Not AI**, record the procurement owner and evidence at every handoff; an arrow is not evidence that the preceding stage passed.

Evidence and control graph

The **buy an outcome, not ai — evidence network** shows the relationships that must remain inspectable when the workflow runs. In *The AI Buyer’s Field Guide*, this qualitative graph asks the reader to record observed evidence for **outcome baseline established** and route **issuing a request for a generic AI platform** to recovery instead of treating sample numbers as facts.




Figure 1.4: Qualitative evidence network for chapter 1, Buy an Outcome, Not AI, linking the question to Write an outcome statement with baseline, target, users, decisions, and non-goals., the measure outcome baseline established, and the recovery condition issuing a request for a generic AI platform without invented empirical values.

Read the graph: Trace the evidence for Buy an Outcome, Not AI toward the decision. For Buy an Outcome, Not AI, empty or weak evidence does not count as zero risk; it routes the buying decision to verification, escalation, or stop. Use this **Buy an Outcome, Not AI** graph beside **Write an outcome statement with baseline, target, users, decisions, and non-goals**, not as a substitute for recorded cases.

Worked composite case file — illustrative, not empirical evidence

This **Buy an Outcome, Not AI** record turns the scenario into an inspectable rehearsal in an AI procurement and contract decision. Its numbers are worksheet examples for Chapter 1, not observed results or a prediction of benefit.

Record field	Filled illustrative entry
Situation	A public-service team reframes a chatbot purchase as reducing unresolved routine inquiries without blocking human access.
People affected	buyers, users, affected data subjects, and vendors
Decision	What business condition should change if the purchase succeeds?
Artifact excerpt	Write an outcome statement with baseline, target, users, decisions, and non-goals.
Accountable authority	The procurement owner with legal and security review.
Measure and unit	outcome baseline established — one dated observation for each eligible case.
Example baseline	7 observations in the baseline window.
Example target / decision rule	a documented improvement plus no severe unresolved exception; move only contextual; the procurement owner must define better before collection.
Named exception	issuing a request for a generic AI platform.
Safe default	Make no commitment and preserve the unresolved requirement.

Scenario trace

1. **Preserve the starting state.** For **Buy an Outcome, Not AI**, record the original request, governing rule, available evidence, and the last safe result before changing anything.
2. **Make the chapter decision.** Answer **What business condition should change if the purchase succeeds?** and record the Chapter 1 assumption that carries the most consequence for buyers, users, affected data subjects, and vendors.
3. **Use the artifact.** Complete **Write an outcome statement with baseline, target, users, decisions, and non-goals** with a normal case and a counter case; in this chapter, do not infer missing facts merely to finish the form.

4. **Observe the decision rule.** Review **outcome baseline established** at each evaluation gate, contract change, and renewal review; the **Buy an Outcome, Not AI** record reports exclusions and review effort beside the headline value.
5. **Contain and recover.** If **issuing a request for a generic AI platform**, make no commitment and preserve the unresolved requirement. Preserve the Chapter 1 failed case, assign a deadline, and require the procurement owner with legal and security review to authorize any restart.

Artifact handoff

Start **Buy an Outcome, Not AI** from the book-level thesis and boundary. Within *The AI Buyer's Field Guide*, attach the current artifact and exception record, then hand both to Chapter 2, **Classify Consequence and Exposure**, which asks: How much diligence does this use case deserve?. That recipient may revise the **Buy an Outcome, Not AI** assumption, but must preserve the evidence that made this book's earlier decision reasonable.

Evidence application

Reference [1], UK Government, *Guidelines for AI procurement*, provides a relevant boundary or design input for **Buy an Outcome, Not AI**. It does **not** prove that the Chapter 1 worked values, local method, or target will succeed. The **Buy an Outcome, Not AI** reviewer must link this book's exact claim to the exact source passage and leave unsupported synthesis marked UNVERIFIED.

Transfer test

- Re-run **Write an outcome statement with baseline, target, users, decisions, and non-goals** with one routine Chapter 1 case that should pass.
- Re-run the **Buy an Outcome, Not AI** artifact with missing evidence and keep the safe default active.
- Re-run it where **issuing a request for a generic AI platform** appears plausible but not certain.
- Ask a reviewer who did not design this chapter's method to decide whether the retained record supports continuing, narrowing, escalating, or stopping.

Measurement and recovery decision

For **Buy an Outcome, Not AI**, define **outcome baseline established** as one dated observation for each eligible case. In this **Buy an Outcome,**

Not AI worksheet, the Chapter 1 worked baseline is 7 observations in the baseline window and the illustrative target is a documented improvement plus no severe unresolved exception; both must be replaced with evidence from an AI procurement and contract decision. Review the Chapter 1 result at each evaluation gate, contract change, and renewal review, preserve excluded and high-consequence cases, and interpret movement as better only when it is contextual; the procurement owner must define better before collection.

The named failure is **issuing a request for a generic AI platform**. Containment means: make no commitment and preserve the unresolved requirement. Recovery is complete only after the procurement owner with legal and security review can reconstruct the failed case, explain which assumption changed, and authorize a new test against **outcome baseline established**.

Practice lab

Choose one live but reversible **Buy an Outcome, Not AI** example from an AI procurement and contract decision. Create the following:

- a one-sentence answer to **What business condition should change if the purchase succeeds?**;
- a completed first draft of **Write an outcome statement with baseline, target, users, decisions, and non-goals**;
- one routine case, one missing-evidence case, and one case shaped by **issuing a request for a generic AI platform**;
- a locally measured baseline for **outcome baseline established** using the unit one dated observation for each eligible case;
- a stop threshold, the procurement owner with legal and security review as owner, and the safe default make no commitment and preserve the unresolved requirement;
- vendor evidence packet that will be handed to the next chapter.

Ask a colleague to challenge **Write an outcome statement with baseline, target, users, decisions, and non-goals**. They should identify one ambiguous field, one person among buyers, users, affected data subjects, and vendors whose interests are missing, and one reason **outcome baseline established** could reward the wrong behavior. Preserve that challenge in the chapter handoff before revising.

Limits and judgment

Buy an Outcome, Not AI supplies a locally testable method, not proof of causal effectiveness or a universal rule. This Chapter 1 method remains inside the book boundary: This book is general information, not legal advice, a model contract, or certification of any supplier or compliance regime. Readers applying **Buy an Outcome, Not AI** for buyers, users, affected data

subjects, and vendors must verify current law, contract, policy, safeguarding, accessibility, security, and professional duties before consequential use. A negative **outcome baseline established** result may mean the question was framed badly, the review cost is too high, or make no commitment and preserve the unresolved requirement is the responsible decision.

Takeaways

- A procurement starts with a measurable use case, affected people, and failure cost rather than a technology label.
- The chapter artifact is: write an outcome statement with baseline, target, users, decisions, and non-goals.
- The primary measure is **outcome baseline established**, paired with review cost and failure severity.
- Watch especially for **issuing a request for a generic AI platform**.
- In an AI procurement and contract decision, expand scope only after the procurement owner with legal and security review can show the evidence and exercise the recovery path.

Reflection

What evidence would let a skeptical reviewer answer **What business condition should change if the purchase succeeds?** differently from the procurement owner with legal and security review? Which case involving buyers, users, affected data subjects, and vendors would **outcome baseline established** miss, and would that omission require continuing, narrowing, escalating, or make no commitment and preserve the unresolved requirement?