

Small Teams, Smart Systems

A Practical Guide to AI Workflows for Owner-Led Businesses

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 Start With Friction, Not Fashion	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 Map the Work as It Happens	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	28
2.6	Practice lab	29
2.7	Limits and judgment	29
2.8	Takeaways	30
2.9	Reflection	30
3	Choose the Smallest Valuable Boundary	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	40
3.6	Practice lab	41
3.7	Limits and judgment	41
3.8	Takeaways	42
3.9	Reflection	42
4	Build a Minimum Useful Data Set	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	52
4.6	Practice lab	53
4.7	Limits and judgment	53
4.8	Takeaways	54
4.9	Reflection	54

5	Design the Human Decision Point	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	64
5.6	Practice lab	65
5.7	Limits and judgment	65
5.8	Takeaways	66
5.9	Reflection	66
6	Prototype With Realistic Cases	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	76
6.6	Practice lab	77
6.7	Limits and judgment	77
6.8	Takeaways	78
6.9	Reflection	78
	Midpoint visual synthesis	79
7	Measure Value and Failure Together	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
7.4	Worked value-and-failure calculation	84
7.4.1	Chapter-specific operating protocol	84
7.4.2	Operating workflow	87
7.4.3	Evidence and control graph	88
7.5	Worked composite case file — illustrative, not empirical evidence	89
7.5.1	Scenario trace	90
7.5.2	Artifact handoff	91
7.5.3	Evidence application	91
7.5.4	Transfer test	91
7.6	Measurement and recovery decision	91
7.7	Practice lab	92
7.8	Limits and judgment	92
7.9	Takeaways	93
7.10	Reflection	93
8	Handle Exceptions as First-Class Work	95
8.1	In plain terms	95
8.1.1	Chapter contract	95
8.2	The decision at stake	95
8.3	Working model	96
8.3.1	See the chapter in practice	96
8.3.2	Chapter-specific operating protocol	98
8.3.3	Operating workflow	100
8.3.4	Evidence and control graph	101
8.4	Worked composite case file — illustrative, not empirical evidence	102
8.4.1	Scenario trace	103
8.4.2	Artifact handoff	104
8.4.3	Evidence application	104
8.4.4	Transfer test	104
8.5	Measurement and recovery decision	104
8.6	Practice lab	105
8.7	Limits and judgment	105
8.8	Takeaways	106
8.9	Reflection	106
9	Select Tools After the Workflow	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	Operate the System Every Week	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	Scale by Reusing Controls, Not Prose	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 The Ninety-Day Operating Plan 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

Start With Friction, Not Fashion

In plain terms

A useful first project is visible, frequent, and reversible. In an owner-led service operation, the practical question is **Which repeated problem is expensive enough to deserve attention?** For owners and operations leads in service businesses with two to fifty people, the answer counts only when **Build a friction inventory and rank items by hours, errors, delay, and customer impact** produces reviewable evidence for **recovered response time** without normalizing **buying a tool before naming the operational loss**.

Chapter contract

DECIDE

Which repeated problem is expensive enough to deserve attention?

BUILD

Build a friction inventory and rank items by hours, errors, delay, and customer impact.

MEASURE

recovered response time

STOP OR REDESIGN WHEN

buying a tool before naming the operational loss

The decision at stake

For **Start With Friction, Not Fashion**, the business owner must decide **which repeated problem is expensive enough to deserve attention?** before the service work crosses from discussion into consequence for customers and frontline staff. The minimum decision record is the artifact **Build a friction inventory and rank items by hours, errors, delay, and customer impact**, the observable **recovered response time**, and the explicit hold condition **buying a tool before naming the operational loss**. If any Chapter 1 field is missing, the safe default is to return the case to the manual service path.

The bounded **Start With Friction, Not Fashion** decision advances the book's promise to turn one recurring business bottleneck at a time into a measured, reviewable AI-assisted workflow. It does not endorse a product or

predetermine the Chapter 1 answer; it creates a record that another person can challenge at the end of each operating week.

Working model

Read the **Start With Friction, Not Fashion** model as four linked fields. **Outcome:** A useful first project is visible, frequent, and reversible. **Boundary:** the design must make **buying a tool before naming the operational loss** visible rather than convenient. **Evidence:** interpret **recovered response time** together with review effort and excluded cases. **Authority:** the business owner may continue, narrow, escalate, or stop the service work. For **Start With Friction, Not Fashion**, Reference [1] from U.S. Small Business Administration, *AI for small business*, informs a Chapter 1 boundary; it is not evidence that the local method or illustrative target will work.

See the chapter in practice

The **start with friction, not fashion in practice** maps *Small Teams, Smart Systems* Chapter 1 into a field scene before its abstract models begin. In one frame, it holds the question **Which repeated problem is expensive enough to deserve attention?**, the artifact **Build a friction inventory and rank items by hours, errors, delay, and customer impact.**, the proof signal **recovered response time**, and the warning **buying a tool before naming the operational loss**.




Figure 1.1: Chapter 1 field scene for Start With Friction, Not Fashion: a boundary view in the book’s small team setting.

Read the scene: For Small Teams, Smart Systems, Chapter 1, begin with the central small team scene and scan DECIDE, MAKE, CHECK, and WATCH. Treat the Start With Friction, Not Fashion scene as conceptual, then use its cases and evidence records to test whether recovered response time actually improves. For **Start With Friction, Not Fashion**, this explanatory plate is a reasoning prompt—not observed evidence or a universal operating prescription.

Visual field note — Start With Friction, Not Fashion:

- 1. Locate the tension.** In this composite situation, A repair company loses weekend inquiries between voicemail and Monday follow-up.
- 2. Name the choice.** Answer the chapter question: **Which repeated problem is expensive enough to deserve attention?**
- 3. Produce the record.** Use **Build a friction inventory and rank items by hours, errors, delay, and customer impact.**, then retain the business owner, assumptions, and case evidence needed to review it.
- 4. Test and return.** Track **recovered response time**. If **buying a tool before naming the operational loss**, return to DECIDE and revise the pilot 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 Which repeated problem is expensive enough to deserve attention?	Chapter 1 owner, input, output, exception, and timestamp.	Hold Start With Friction, Not Fashion if its record cannot be reconstructed.
2. Contrast	Apply Build a friction inventory and rank items by hours, errors, delay, and customer impact to one ordinary case.	Chapter 1 owner, input, output, exception, and timestamp.	Hold Start With Friction, Not Fashion if its record cannot be reconstructed.
3. Exclude	Run the same artifact against a countercase shaped by buying a tool before naming the operational loss .	Chapter 1 owner, input, output, exception, and timestamp.	Hold Start With Friction, Not Fashion if its record cannot be reconstructed.
4. Assign	Name the business owner as the person who may continue, narrow, or stop.	Chapter 1 owner, input, output, exception, and timestamp.	Hold Start With Friction, Not Fashion if its record cannot be reconstructed.
5. Challenge	Record recovered response time and the review effort for the same case.	Chapter 1 owner, input, output, exception, and timestamp.	Hold Start With Friction, Not Fashion if its record cannot be reconstructed.

Phase	Reader action	Evidence retained	Hold condition
6. Publish	Release the Start With Friction, Not Fashion artifact only when its evidence answers this chapter's question.	Chapter 1 owner, input, output, exception, and timestamp.	Hold Start With Friction, Not Fashion if its record cannot be reconstructed.

This sequence is specific to **Start With Friction, Not Fashion**. For that Chapter 1 decision, it is not permission to automate; it is a way for owners and operations leads in service businesses with two to fifty people to make the reasoning inspectable at the end of each operating week.

Build a friction inventory and rank items by hours, errors, delay, and customer impact. This is the chapter's core production artifact. The **Start With Friction, Not Fashion** 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 **Start With Friction, Not Fashion**, it translates explanation into an inspectable decision aid: the reader can see the decision, artifact, evidence, and exact condition that pauses the service work.




Figure 1.2: Chapter contract connecting the decision, artifact, measure, and stop condition for Start With Friction, Not Fashion.

How to use this page: Move across the four Start With Friction, Not Fashion cards; compare the decision with recovered response time and stop before buying a tool before naming the operational loss. The chapter-specific measure is **recovered response time**; the principal warning is **buying a tool before naming the operational loss**.

Operating workflow

The **start with friction, not fashion workflow** makes the chapter executable. For **Start With Friction, Not Fashion**, the five stages carry **Build a friction inventory and rank items by hours, errors, delay, and customer impact.** from the question **Which repeated problem is expensive enough to deserve attention?** to a test against **recovered response time**, with **buying a tool before naming the operational loss** as an explicit return or stop signal.




Figure 1.3: Five-stage workflow for chapter 1, Start With Friction, Not Fashion: frame the decision, set the pilot boundary, build Build a friction inventory and rank items by hours, errors, delay, and customer impact., test against recovered response time, then continue, narrow, or stop before buying a tool before naming the operational loss.

Read the workflow: Move from FRAME to DECIDE for Start With Friction, Not Fashion. The lower return path is mandatory when evidence is missing or the stop signal—buying a tool before naming the operational loss—appears. For **Start With Friction, Not Fashion**, record the business owner and evidence at every handoff; an arrow is not evidence that the preceding stage passed.

Evidence and control graph

The **start with friction, not fashion — evidence network** shows the relationships that must remain inspectable when the workflow runs. In *Small Teams, Smart Systems*, this qualitative graph asks the reader to record observed evidence for **recovered response time** and route **buying a tool before naming the operational loss** to recovery instead of treating sample numbers as facts.




Figure 1.4: Qualitative evidence network for chapter 1, Start With Friction, Not Fashion, linking the question to Build a friction inventory and rank items by hours, errors, delay, and customer impact., the measure recovered response time, and the recovery condition buying a tool before naming the operational loss without invented empirical values.

Read the graph: Trace the evidence for Start With Friction, Not Fashion toward the decision. For Start With Friction, Not Fashion, empty or weak evidence does not count as zero risk; it routes the service work to verification, escalation, or stop. Use this **Start With Friction, Not Fashion** graph beside **Build a friction inventory and rank items by hours, errors, delay, and customer impact**, not as a substitute for recorded cases.

Worked composite case file — illustrative, not empirical evidence

This **Start With Friction, Not Fashion** record turns the scenario into an inspectable rehearsal in an owner-led service operation. Its numbers are worksheet examples for Chapter 1, not observed results or a prediction of benefit.

Record field	Filled illustrative entry
Situation	A repair company loses weekend inquiries between voicemail and Monday follow-up.
People affected	customers and frontline staff
Decision	Which repeated problem is expensive enough to deserve attention?
Artifact excerpt	Build a friction inventory and rank items by hours, errors, delay, and customer impact.
Accountable authority	The business owner.
Measure and unit	recovered response time — eligible cases meeting the defined condition per 100 cases.
Example baseline	63%.
Example target / decision rule	at least 79%; move only higher, while guardrails remain stable.
Named exception	buying a tool before naming the operational loss.
Safe default	Return the case to the manual service path.

Scenario trace

1. **Preserve the starting state.** For **Start With Friction, Not Fashion**, record the original request, governing rule, available evidence, and the last safe result before changing anything.
2. **Make the chapter decision.** Answer **Which repeated problem is expensive enough to deserve attention?** and record the Chapter 1 assumption that carries the most consequence for customers and frontline staff.
3. **Use the artifact.** Complete **Build a friction inventory and rank items by hours, errors, delay, and customer impact** with a normal case and a countercase; in this chapter, do not infer missing facts merely to finish the form.
4. **Observe the decision rule.** Review **recovered response time** at the end of each operating week; the **Start With Friction, Not Fashion** record reports exclusions and review effort beside the headline value.
5. **Contain and recover.** If **buying a tool before naming the operational loss**, return the case to the manual service path.

Preserve the Chapter 1 failed case, assign a deadline, and require the business owner to authorize any restart.

Artifact handoff

Start **Start With Friction, Not Fashion** from the book-level thesis and boundary. Within *Small Teams, Smart Systems*, attach the current artifact and exception record, then hand both to Chapter 2, **Map the Work as It Happens**, which asks: What actually occurs between a request and a completed outcome?. That recipient may revise the **Start With Friction, Not Fashion** assumption, but must preserve the evidence that made this book's earlier decision reasonable.

Evidence application

Reference [1], U.S. Small Business Administration, *AI for small business*, provides a relevant boundary or design input for **Start With Friction, Not Fashion**. It does **not** prove that the Chapter 1 worked values, local method, or target will succeed. The **Start With Friction, Not Fashion** reviewer must link this book's exact claim to the exact source passage and leave unsupported synthesis marked UNVERIFIED.

Transfer test

- Re-run **Build a friction inventory and rank items by hours, errors, delay, and customer impact** with one routine Chapter 1 case that should pass.
- Re-run the **Start With Friction, Not Fashion** artifact with missing evidence and keep the safe default active.
- Re-run it where **buying a tool before naming the operational loss** 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 **Start With Friction, Not Fashion**, define **recovered response time** as eligible cases meeting the defined condition per 100 cases. In this **Start With Friction, Not Fashion** worksheet, the Chapter 1 worked baseline is 63% and the illustrative target is at least 79%; both must be replaced with evidence from an owner-led service operation. Review the Chapter 1 result at the end of each operating week, preserve excluded and high-consequence cases, and interpret movement as better only when it is higher, while guardrails remain stable.

The named failure is **buying a tool before naming the operational loss**. Containment means: return the case to the manual service path. Recovery is complete only after the business owner can reconstruct the failed case, explain which assumption changed, and authorize a new test against **recovered response time**.

Practice lab

Choose one live but reversible **Start With Friction, Not Fashion** example from an owner-led service operation. Create the following:

- a one-sentence answer to **Which repeated problem is expensive enough to deserve attention?**;
- a completed first draft of **Build a friction inventory and rank items by hours, errors, delay, and customer impact**;
- one routine case, one missing-evidence case, and one case shaped by **buying a tool before naming the operational loss**;
- a locally measured baseline for **recovered response time** using the uneligible cases meeting the defined condition per 100 cases;
- a stop threshold, the business owner as owner, and the safe default return the case to the manual service path;
- operating evidence packet that will be handed to the next chapter.

Ask a colleague to challenge **Build a friction inventory and rank items by hours, errors, delay, and customer impact**. They should identify one ambiguous field, one person among customers and frontline staff whose interests are missing, and one reason **recovered response time** could reward the wrong behavior. Preserve that challenge in the chapter handoff before revising.

Limits and judgment

Start With Friction, Not Fashion 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 does not promise a revenue result, replace legal or accounting advice, or certify any vendor. Readers applying **Start With Friction, Not Fashion** for customers and frontline staff must verify current law, contract, policy, safeguarding, accessibility, security, and professional duties before consequential use. A negative **recovered response time** result may mean the question was framed badly, the review cost is too high, or return the case to the manual service path is the responsible decision.

Takeaways

- A useful first project is visible, frequent, and reversible.
- The chapter artifact is: build a friction inventory and rank items by hours, errors, delay, and customer impact.
- The primary measure is **recovered response time**, paired with review cost and failure severity.
- Watch especially for **buying a tool before naming the operational loss**.
- In an owner-led service operation, expand scope only after the business owner can show the evidence and exercise the recovery path.

Reflection

What evidence would let a skeptical reviewer answer **Which repeated problem is expensive enough to deserve attention?** differently from the business owner? Which case involving customers and frontline staff would **recovered response time** miss, and would that omission require continuing, narrowing, escalating, or return the case to the manual service path?