

AI BUSINESS OPERATING SYSTEM

Reference Cards & Appendix Tools

Ten practical reference cards reproduced from the book appendices.

Appendices

Ten reference cards. Each is a page you will use more often than the chapter it came from. Reproduce them, print them, and put them where the work happens.

Appendix A — The Implementation Spine

Thirteen steps, three phases, a gate at every step. Every automation in this book follows it in this order. (Chapter 3; Figure 3.4.)

#	Step	Phase	Gate — must hold before you advance	Fullest treatment
1	Audit the work	DECIDE	A ranked queue exists with at least one task scoring 3/4 or 4/4	Ch 1
2	Choose the stack	DECIDE	All four Stack Ladder rungs named; Store data rules written	Ch 2
3	Learn to brief	DECIDE	A BRIEF-format prompt has produced acceptable output twice in a row	Ch 3
4	Set the target	DECIDE	A KPI entry with a stated formula exists before building continues	Ch 7
5	Gather the sources	BUILD	At least one source has passed ARS scoring and is logged in the Approved Documents Register	Ch 5
6	Verify before you trust	BUILD	Every source in use has a filled Tier and Expiry field	Ch 6
7	Assemble the workflow	BUILD	Workflow is mapped stage by stage before any AI step is wired in	Ch 8

8	Add the human gate	BUILD	A named human checkpoint exists with a named individual owner	Ch 13
9	Test on real cases	BUILD	A stated minimum of real historical cases tested, with recorded pass rate	Ch 9
10	Route the exceptions	SUSTAIN	A named escalation path exists with a stated response time	Ch 16
11	Write the procedure	SUSTAIN	A Role Handbook entry exists with a named owner and stated authority boundary	Ch 17
12	Review on schedule	SUSTAIN	A specific review date is calendared	Ch 19
13	Retire or replace	SUSTAIN	A written exit criterion exists before the automation is considered complete	Ch 20

Keep a gate log: one row per automation — which step it is at, who confirmed each gate, and the date. “Is this finished?” must have a written answer, not a remembered one.

Appendix B — The DRIP Test

Four criteria, one point each, applied to every recurring task. Used once, at the start, to decide whether a task is a candidate for automation and what to do with it. (Chapter 1; Figure 1.1.)

	Criterion	Scores 1 if
D	Defined	You can write the task's start state and end state in one paragraph, without a meeting
R	Repeatable	It recurs on a fixed schedule, in the same form, rather than surfacing once
I	Independent	It can be completed with information already at hand, without another person's live judgment call in the middle
P	Predictable	The range of acceptable outputs is narrow enough that a wrong answer would be obviously wrong

Score	Decision
4/4	Automate now

3/4	Automate with a checkpoint — automate the mechanical part, insert a human review where judgment remains
2/4	Redesign the task first — usually Defined or Independent is failing, and no tool fixes an undefined task
0–1/4	Leave it alone

Capture before you score. One full working week, every task that repeats, however small. Memory produces eight to twelve tasks; a week of writing produces thirty to fifty. The difference is where the leverage is.

Rank after you score. First by DRIP score, then within each band by minutes reclaimed per week (frequency × time per instance). The result is the ranked queue — Worksheet 1.1.

The audit does not use AI. Scoring is a judgment about your business, made by the person who does the task or manages the person who does. The task owner signs off the score. Re-run the capture at least twice a year.

Appendix C — The Stack Ladder

Four rungs, always climbed in the same order. (Chapter 2; Figures 2.1–2.3.)

Rung	Name	What happens here	The usual failure
1	Capture	Information enters the business: forms, emails, scans, voice notes, photographs, uploads, spreadsheets	Skipped: nothing downstream can work with information that never entered in a usable form
2	Connect	Captured information is routed to the systems and people that need it, without a human copying it between them	Skipped: each tool captures and processes beautifully, then strands the result
3	Compute	AI does the processing: drafting, summarizing, classifying, extracting, analyzing	Over-invested: this is the rung every demo shows
4	Store	Verified output becomes a permanent business record rather than a conversation that evaporates	Never built voluntarily; the single biggest predictor of whether adoption compounds or resets

A human checkpoint sits between Compute and Store. Nothing enters permanent record without someone confirming it is correct. It is never the first thing to cut when moving faster.

One shared Connect layer per function, not one connection per source system. Five point-to-point connections fail silently; one routing layer can be checked in one place.

Store is governed under data protection from day one. Three rules regardless of industry: tell the data subject at Capture what the data will be used for; never keep personal data longer than the stated purpose requires, and build the deletion step into the stack; treat sensitive personal data — health, financial, biometric — as needing a higher bar of consent and access at every rung.

The Business Brain lives at the Store rung. A working Store rung is the one prerequisite Chapter 5 assumes.

Appendix D — The BRIEF Method

Five parts plus a mandatory RULES block. The structure that turns an unreliable prompt into a reliable one. (Chapter 3; Figures 3.1–3.3; Worksheet 3.1.)

Element	What goes in it
Background	The context the AI cannot infer: what the business does, who the output is for, what has already been tried
Role	Who the AI is acting as, and at what level of authority — chosen deliberately, not by accident
Instructions	The specific task, as actions rather than goals, broken into steps if it has more than one
Examples	At least one worked example of the output you want, in the exact format you want it. Removes more failures than any other element
Format	The exact structure the output must take
RULES	Hard constraints the AI must never violate: never invent a figure it was not given; never make a decision that requires human judgment; never proceed if required information is missing — stop instead of guessing. The longest section, deliberately

The loop. Write → test against real cases → evaluate → if it fails, patch one line (usually in RULES) → test again. Stop when it reliably produces correct output, not when it produces output that looks correct once.

The gate. Acceptable output twice in a row on real material. Test at least three real cases before calling a prompt complete.

When it fails, diagnose in two minutes (Figure 3.3): wrong format → Format or Examples is missing; invented facts → RULES is missing or vague; wrong task done well → Instructions state a goal, not an action; wrong tone or authority → Role was accidental; answer that ignores your situation → Background is thin.

When you adapt a printed prompt, change the bracketed fields. Keep the rules.

Appendix E — ARS Scoring

Authority · Recency · Specificity. The test for whether a source belongs in the Business Brain at all. Eight factors, one point each. (Chapter 5; Figure 5.1; Worksheet 5.1.)

	Factor	Scores 1 if
A1	Named	The source is named and identifiable, not anonymous
A2	Direct	The source has direct authority over the subject — the vendor's own manual, not a forum summary of it
A3	Verifiable	It can be checked against a primary document if challenged
R1	Dated	Its publication or issue date is known

R2	Current	It is within its stated validity window and not superseded
S1	On point	It directly addresses the exact question it will be used to answer
S2	Actionable	It contains specific, usable detail rather than general background
S3	Consistent	It contradicts nothing already in the Register

Score	Status	Handling
7–8	Approved	Enters the Register, no special handling
5–6	Approved with review date	Enters the Register, flagged for periodic re-check
3–4	Provisional	Usable only alongside a human check every time it is cited
0–2	Rejected	Does not enter the Business Brain

A document with an unknown date never scores above 4, however authoritative it looks. Tax law, regulatory guidance and vendor specifications change more often than the people using them assume.

Score with the person who relies on the document, not as a solo desk exercise — especially the Recency factors. AI can draft scoring notes and flag likely Recency problems; the Authority and Specificity judgments need someone with subject knowledge.

ARS scores a source's reliability. It does not score whether using that source for a given purpose is appropriate — that remains a judgment about fit.

Appendix F — The OACI Ownership Model

Exactly one of four roles, for every relevant person, for every documented procedure, workflow and system. (Chapter 17; Figure 17.3; Worksheet 17.1.)

Role	Meaning	Mark on the matrix
O — Owner	The single person accountable for the procedure's accuracy and currency, and the only role who may approve a change to it	Filled circle
A — Approver	Must sign off before a proposed change takes effect — distinct from the Owner, so no one approves their own unchecked changes	Filled square
C — Contributor	May use the procedure and suggest	Outlined triangle

	changes; cannot approve or make them unilaterally	
I — Informed	Must be notified when the procedure changes; no approval rights and no obligation to act	Small outlined circle

The matrix is procedures against people, one shape per cell, with a printed key. It must stay legible in pure black and white.

Review it against actual staff turnover. A Role Handbook that still names someone who left months ago is a continuity document that has stopped working. Ownership is reassigned formally through this matrix, not by assumption.

Appendix G — The Thirteen Artefacts of the Business Brain

Built one per chapter, Chapter 5 to Chapter 17, in this order. Exactly thirteen. Chapters 18–20 draw on all of them and add none.

#	Artefact	Ch	Worksheet	Spine step it serves
1	Approved Documents Register	5	5.1	5 · Gather the sources
2	Source Register and Findings Log	6	6.1	6 · Verify before you trust
3	KPI Definition Sheet	7	7.1	4 · Set the target
4	Content Pillars and Voice Guide	8	8.1	7 · Assemble the workflow
5	Knowledge Asset and Derivatives Register	9	9.1	5 · Gather the sources
6	Social Operating System	10	10.1	7 · Assemble the workflow
7	Recording Plan and Media Asset Register	11	11.1	8 · Add the human gate
8	Visual Specification and Template Backlog	12	12.1	7 · Assemble the workflow
9	The Qualification Standard	13	13.1	8 · Add the human gate
10	Messaging Library	14	14.1	7 · Assemble the workflow
11	Commitment Log	15	15.1	9 · Test on real cases
12	Escalation Standard	16	16.1	10 · Route the exceptions

13	Role Handbook	17	17.1	9 · Test on real cases
----	---------------	----	------	------------------------

Every artefact is a document first and a tool's data second. Written, scored, owned — never a folder of unverified material. Each carries a named owner and a review date.

Every artefact passes the same review on schedule (Figure 19.2): Is it still current against what the business actually does? If not, update before trusting it again. Is its named owner still with the business and still the right person? If not, reassign through the OACI matrix before proceeding.

The question the complete set answers (Chapter 19): what would this business lose if one specific person were unavailable tomorrow?

Appendix H — The AI Boundary

Every chapter draws the same line in its own function. Collected here: what AI may do, and what a person must decide. (The AI Boundary section of Chapters 1–20.)

Ch	AI may	A person must
1	Surface candidate tasks you might not think to list	Score and rank them. The audit itself does not use AI
2	Draft the technical specification for a Connect-layer integration	Decide which rung a tool occupies and whether a purchase fills a real gap
3	Draft a prompt, diagnose a failing one, suggest what RULES might need	Certify every gate: enough real cases tested, the right checkpoint owner, ready to advance
4	Categorize email and draft replies	Decide what a customer is owed, what tone a relationship calls for, whether an unusual request is honored
5	Draft ARS scoring notes and flag Recency issues	Make the Authority and Specificity judgments; confirm a source is current
6	Locate, summarize and trace a finding toward its source	Certify that a source is Tier 1
7	Explain, in plain language, why an approved number moved	Choose what to measure, set the target, review every explanation before it is shown
8	Draft content against a defined voice and pillar set	Build and maintain the prohibited-claims list from direct regulatory knowledge
9	Draft derivative formats from an approved parent asset	Decide which topics deserve to become a Knowledge Asset
10	Draft platform-adapted content; auto-reply to general inquiries	Be the first response to any complaint or crisis-level incident
11	Transcribe; generate highlight candidates	Confirm an assembled sequence still represents what was actually said
12	Draft new assets against an approved	Decide what belongs in the minimum

	template	specification
13	Apply a written disqualifier list and ICP criteria consistently	Decide what is on the disqualifier list
14	Draft stage messages; check contacts against the suppression matrix mechanically	Define the engagement threshold in writing before anything is judged “disengaged”
15	Draft commitment-seeking follow-ups; flag stale deals	Insist on a specific date and name before logging anything as a commitment
16	Check contacts against the approved FAQ library; draft auto-replies	Decide whether a never-automate category applies — conservatively, literally, erring toward a person
17	Transcribe a narration; draft a first-pass procedure	Test the procedure by having someone who did not know it follow it exactly
18	Operate within defined authority limits; flag its own low-confidence decisions	Set every authority limit; make every change of level a deliberate decision, never an inferred one
19	Compile the weekly view from already-logged data	Judge whether an artefact’s content is still accurate
20	Draft a data-portability request; compare export structures	Perform the actual export test against real data

Two things never move to the machine: the decision that requires judgment about your specific business, and the confirmation that something is true.

Appendix I — Governance Touchpoints

Where the book’s automations meet an obligation. Each is stated in the Governance section of the chapter named; none is legal advice, and each ends the same way — take advice on your own situation. (Chapters 1–20.)

Obligation	Where it bites	Ch
Ownership of the audit sits with whoever runs the function, not a central “AI team”; nothing is built without the task owner’s sign-off	The Leverage Audit	1
Personal data at Store needs a lawful basis, a stated purpose and a retention period — India’s Digital Personal Data Protection Act, 2023	The Stack Ladder	2
Every automation has a gate log: step, who confirmed, date	The Implementation Spine	3
Retained email content needs a stated purpose, retention period, restricted access; complaint correspondence is	Inbox triage	4

never used to train an external model without consent		
In regulated professions, an AI-drafted response citing an Approved Document is not professional advice until the required professional signs it off	Approved Documents Register	5
Findings about named competitors or individuals are logged factually, sourced and dated — never speculation presented as fact	Source Register	6
Dashboards show aggregates; individual-level drill-down needs a logged, restricted access step	KPI dashboard	7
Claims about safety, efficacy, pricing guarantees or comparative performance route to compliance review regardless of how routine the piece is	Content engine	8
A Knowledge Asset touching regulated claims clears claims review before any derivative is produced — a wrong claim in ten formats is a wrong claim ten times faster	Knowledge Assets	9
Commercial messaging by WhatsApp, SMS or similar falls under TRAI's framework: recipient consent, the Do Not Disturb registry, registered channels — whoever drafted the message	Social Operating System	10
Documented consent for recording and for its intended use, before production; spoken claims route through claims review	Recording and release	11
The minimum specification and every approved template sit in the Approved Documents Register, ARS-scored, with the brand owner named	Visual Specification	12
Disqualifiers reflect genuine business constraints, never protected characteristics or a proxy for them	Qualification Standard	13
Suppression logic governs marketing and relationship messaging only — never a customer's inbound request for support	Lifecycle messaging	14
Stakeholder positions in a deal record are personal data and do not circulate beyond the deal team	Commitment Log	15
Any request to access, correct or delete personal data is on the never-automate	Escalation Standard	16

list and needs a documented grievance process with a named grievance officer		
The Role Handbook is reviewed against actual turnover so Owner and Approver assignments stay current	Role Handbook	17
A Level 3 agent's actions are logged in a form a human can audit afterwards	Agents	18
The completed consequence audit lists the business's single points of failure and is stored with continuity-record access discipline	Consequence audit	19
Vendor contracts holding customer personal data carry a binding data-return or data-deletion clause at termination	Tool inventory	20

Appendix J — Automation Levels, Tool Lifecycle and Exit

Three levels of automation (Chapter 18; Figure 18.2):

Level	Name	Who acts
1	Assisted	A person performs the task; AI drafts, suggests or checks; the person decides every time
2	Supervised automation	AI executes; a person approves each individual action before it takes effect
3	Autonomous agent	AI executes and acts without per-instance approval, inside authority limits set in advance

Three preconditions before any Level 3 authority (Worksheet 18.1), each answered with evidence, not impression: Is every action fully reversible? Has it run at Level 2 with a logged error rate? Is the worst-case cost acceptable? Then an authority table — may act without escalation · must escalate · never permitted — and a named review owner with a review frequency. An agent never sets its own limits, promotes itself to Level 3, or widens its own threshold on a good track record.

The tool lifecycle (Chapter 20; Figure 20.2): Evaluate (including an export-path check) → Adopt → Operate → Reassess, on a fixed schedule → Exit or Renew → back to Evaluate for whatever replaces it. No tool sits outside the cycle. “We’ve always used this one” is not a reason to skip a reassessment.

Before any renewal (Figure 20.3): Does a verified, tested export path for our accumulated data exist? No → close that gap before renewing, whatever else is true. Yes → does the tool still meet current needs at acceptable cost? Yes → renew. No → begin the exit while the export path is already confirmed.

What is portable (Figure 20.1): documented knowledge — the thirteen artefacts — is portable by design; configured workflows and automations are migratable with effort; data trapped in a vendor's format with no verified export path is the risk zone.

The inventory (Worksheet 20.1): thirteen rows, one per artefact, in chapter order — which tool holds it, whether its export path has been tested, and by whom, on what date. An export path is marked verified only after someone has performed it, successfully, against real data, with the date recorded. A vendor's advertised feature is unconfirmed until then.