

THE ONE-PERSON AI DEPARTMENT TOOLKIT

Companion Resources for *The One-Person AI Department*

James Inglis

The One-Person Business Playbooks · Playbook 04 · Capability

The One-Person AI Department Toolkit

Companion Resources for The One-Person AI Department

James Inglis

The One-Person Business Playbooks · Playbook 04 · Capability

How to Use This Toolkit

This toolkit is designed to help you turn the ideas in The One-Person AI Department into working systems. Do not feel obligated to complete every page. Start with the decision or workflow in front of you, use the relevant template, and improve it as you learn.

A practical starting sequence is:

1. Audit recurring work for AI fit.
2. Define the few AI roles you actually need.
3. Build your Company Brain and core business context.
4. Turn repeated work into documented workflows.
5. Increase automation only after the workflow is reliable.
6. Measure time saved, quality, errors, and owner involvement.

The tools are intentionally simple. Their value comes from repeated use, not from completing them perfectly the first time.

Core Templates

The templates that follow are designed to be copied, edited, and reused. Keep them simple. The best template is the one you actually reach for when work begins.

Template 1 — AI Opportunity Audit

Purpose: Identify recurring work where AI can create useful leverage.

Field	What to Enter
Activity	The recurring task or process
Business Function	Executive, marketing, sales, delivery, support, operations, finance, admin, research
Trigger	What causes the work to begin
Frequency	Daily, weekly, monthly, event-driven
Minutes Each Time	Current owner time
Repeatability (1–5)	How similar the process is each time
Business Value (1–5)	How much the outcome matters
Judgment Required (1–5)	How much experienced human judgment is needed
Risk (1–5)	Consequence if the result is wrong
AI Fit (1–5)	How well AI can support the work
Current Tools	Systems already involved
Recommended Bucket	Eliminate / Human / Assist / Automate
Estimated Weekly Time Saved	Conservative estimate
First Experiment	Smallest useful test

Simple Prioritization Rule

Start with work that scores high on frequency, repeatability, value, and AI fit, while remaining low to moderate on risk.

Don't chase the highest theoretical time saving if the workflow is difficult to control.

Template 2 — AI Department Org Chart

Owner / CEO

↓

AI Chief of Staff

- Daily and weekly prioritization
- Meeting preparation
- Commitment tracking
- Decision support

AI Research Analyst

- Market research
- Competitor intelligence

- Customer insight
- Decision research

AI Growth Specialist

- Campaign support
- Content systems
- Prospect research
- Sales follow-up
- Proposal preparation

AI Operations Coordinator

- Project summaries
- SOP creation
- Action tracking
- Status reporting

AI Customer Support Specialist

- Question classification
- Knowledge retrieval
- Draft responses
- Escalation

AI Back-Office Analyst

- Document processing
- Vendor comparisons
- Financial preparation
- Reporting support

Customize It

Delete any role that has no recurring work. Add a specialized role only when the job is real.

Template 3 — AI Role Card

Role Name:

Mission:

One sentence describing why this role exists.

Responsibilities:

- 1.
- 2.
- 3.
- 4.
- 5.

Inputs Required:

-
-

-

Approved Resources:

- Company documents:
- Data sources:
- External sources:
- Tools:

Expected Outputs:

-

-

-

Quality Standard:

Describe what a strong output must contain and what it must avoid.

Authority Level:

0 — Suggest only

1 — Prepare drafts

2 — Ask permission before action

3 — Execute within defined limits

4 — Execute and report exceptions

May Act Without Approval When:

-

-

Must Escalate When:

-

-

-

Known Failure Modes to Watch:

-

-

Version / Last Reviewed:

Template 4 — Company Brain Builder

A. Company

What we do:

Who we serve:

How we make money:

Current priorities:

Important constraints:

What we are not trying to be:

B. Customers

Priority segments:

Typical buying triggers:

Top problems:

Buying criteria:

Common objections:

Decision makers / influencers:

Language customers use:

Strong-fit customer examples:

Poor-fit customer examples:

C. Products and Offers

For each offer:

Name:

Target customer:

Problem solved:

What is included:

What is not included:

Pricing / pricing logic:

Differentiators:

Proof:

Common questions:

D. Brand

Positioning:

Tone:

We should sound:

We should never sound:

Words / phrases we use:

Words / phrases we avoid:

Approved examples:

E. Operations

Major workflow stages:

Approval rules:

Standard policies:

Escalation rules:

Naming / storage conventions:

Recurring reporting cadence:

F. Proof

Case studies:

Testimonials:

Measured outcomes:

Certifications / credentials:

Approved factual claims:

Claims we should not make:

Template 5 — Core Business Context Document

Use this as the concise orientation file shared across AI roles.

1. Company Overview

Who we are, how the business works, and what makes it distinctive.

2. Current Goals

The three to five outcomes that matter most this quarter or year.

3. Target Customers

Who we serve, what matters to them, and how they buy.

4. Offers

What we sell, to whom, at what price or pricing logic, and how the offers relate.

5. Positioning

Why customers should choose us rather than a competitor, substitute, or doing nothing.

6. Brand Voice

How we communicate and examples of language that sounds right.

7. Operating Rules

Important workflow, approval, service, and escalation rules.

8. Proof

Approved evidence supporting major claims.

9. Do Not Assume

List common mistakes an AI system might make about your company.

Examples:

- Don't assume every prospect is a fit.
- Don't assume we discount.
- Don't promise delivery dates unless they appear in the approved system.
- Don't cite a case study without checking that we are authorized to name the customer.

10. Current Priorities

Short-lived context that should be updated frequently.

Last updated:

Next review date:

Template 6 — Universal AI Workflow Builder

Workflow Name:

Purpose:

What business outcome should this workflow produce?

1. Trigger

What starts the workflow?

2. Inputs

What information must be available before work begins?

3. Instructions

What is the objective? What rules and constraints apply?

4. AI Work

Step 1:

Step 2:

Step 3:

Step 4:

5. Human Check

Where does the owner review, decide, approve, or add judgment?

6. Output

What exact deliverable should be produced?

7. Destination

Where does the output go next?

8. Exception Rules

Stop and escalate when:

-
-
-

9. Measures

- Owner minutes required:
- Cycle time:
- Acceptance rate:
- Error rate:
- Business impact:

10. Version

Date created:

Last revised:

Next review:

Template 7 — AI Automation Canvas

Process:

Current Owner Time:

Frequency:

Current Pain Point:

What Can Be Standardized?

What Still Requires Judgment?

Recommended Autonomy Level: 0 / 1 / 2 / 3 / 4

Tools Required:

Data Required:

Permissions Required:

External Actions the System May Take:

External Actions the System May Not Take:

Approval Point:

Exception Conditions:

Logging Method:

Kill Switch:

Success Measure:

Conditions Required Before Increasing Autonomy:

Template 8 — AI Quality-Control Checklist

Use the relevant items rather than blindly checking every box for every output.

Facts

- Names are correct.
- Dates are correct.
- Numbers are verified.
- Quotes are accurate and permitted.
- Sources support material claims.
- Current information is actually current enough for the purpose.

Context

- The output reflects the correct customer, project, or market.
- No outdated company information has been used.
- Assumptions are visible.
- Missing information has not been silently invented.

Commercial Accuracy

- Pricing is approved.
- Scope is accurate.
- Terms are correct.
- Commitments are realistic and authorized.

Brand

- Tone sounds like the company.
- The writing avoids hype and generic filler.
- Claims are supported by approved proof.
- Customer language is used appropriately.

Risk

- Sensitive information is handled appropriately.
- Legal, financial, safety, or policy issues have been escalated where needed.
- The output doesn't exceed the role's authority.

Final Read

- A human has read the final version when the consequence warrants it.

Template 9 — Weekly AI Department Scorecard

Week of:

Complete one copy of the scorecard below for each important workflow you want to track that week.

Field	Weekly Result
Workflow	
Runs	
Owner Time Before	
Owner Time Now	
Time Saved	
Acceptance Rate	
Errors / Exceptions	
Business Impact	
Next Improvement	

Weekly Summary

Most valuable workflow this week:

Biggest failure or surprise:

Context that needs updating:

Workflow ready for greater autonomy:

Workflow that should be simplified:

New process worth testing:

Template 10 — 30-Day Implementation Planner

Week 1 — Audit

- Build work inventory.
- Estimate recurring owner time.
- Score AI opportunities.
- Select five candidates.
- Choose first three workflows.

Week 2 — Foundation

- Define first three AI roles.
- Choose primary AI environment.
- Build Core Business Context.

- Create role context packs.
- Organize approved source documents.

Week 3 — Workflows

- Build executive workflow.
- Build research workflow.
- Build revenue or delivery workflow.
- Run each manually at least three times.
- Record failure modes and revisions.

Week 4 — Automation and Measurement

- Select one stable workflow for automation.
- Set autonomy level.
- Add guardrails and escalation rules.
- Add logging.
- Measure baseline and current performance.
- Hold first AI Operations Review.

End-of-Month Questions

What now happens consistently that did not happen consistently thirty days ago?

What work no longer requires me to start from a blank page?

Where did I save meaningful attention, not merely minutes?

What should the department learn next?

25 Starter AI Workflows

These workflows are intentionally described at the business-process level. Adapt the tools to your stack rather than adapting your business to a tool.

Workflow 1 — Morning CEO Briefing

Trigger: Start of workday.

Inputs: Calendar, open commitments, project deadlines, priority list, notable overnight changes.

AI work: Produce a one-page brief with top three outcomes, meetings requiring preparation, commitments due, decisions waiting, and watch items.

Human role: Confirm priorities and make tradeoffs.

Output: Daily brief.

Best autonomy level: 1 or 2.

Workflow 2 — Weekly Business Review

Trigger: End of week.

Inputs: Sales, projects, finance snapshot, customer issues, commitments, prior weekly plan.

AI work: Summarize progress, stalled items, major changes, decisions needed, and suggested focus for next week.

Human role: Interpret, decide, and set next priorities.

Output: One-page CEO review.

Best autonomy level: 1.

Workflow 3 — Meeting Preparation

Trigger: Important meeting scheduled within 24 hours.

Inputs: Calendar event, attendees, account history, prior notes, current project or opportunity status.

AI work: Prepare background, objective, open commitments, likely discussion points, questions, and sensitive issues.

Human role: Review and decide meeting strategy.

Output: Five-minute meeting brief.

Best autonomy level: 2 or 3 once stable.

Workflow 4 — Meeting to Action

Trigger: Transcript or notes arrive.

Inputs: Transcript, meeting purpose, relevant project context.

AI work: Extract decisions, commitments, actions, owners, dates, unresolved questions, and draft follow-up.

Human role: Validate commitments and approve external message.

Output: Meeting brief plus follow-up.

Best autonomy level: 2.

Workflow 5 — Decision Memo

Trigger: Owner identifies a decision worth structured analysis.

Inputs: Decision question, facts, constraints, relevant research.

AI work: Separate facts and assumptions, lay out options, benefits, risks, reversibility, missing information, and decision criteria.

Human role: Make the decision.

Output: One- to two-page memo.

Best autonomy level: 1.

Workflow 6 — Competitor Intelligence Brief

Trigger: Weekly or monthly schedule.

Inputs: Defined competitor list and approved sources.

AI work: Identify material changes, new offers, messaging changes, leadership moves, partnerships, pricing signals, and possible implications.

Human role: Decide what matters strategically.

Output: Concise competitor brief with sources and dates.

Best autonomy level: 3 for collection; 1 for interpretation.

Workflow 7 — Market Signal Monitor

Trigger: Scheduled scan.

Inputs: Defined market topics, sources, keywords, and company strategy.

AI work: Surface genuinely new developments and group them by opportunity, risk, or neutral change.

Human role: Decide whether any signal deserves action.

Output: Short watch list.

Best autonomy level: 3 or 4 with exception reporting.

Workflow 8 — Customer Interview Analysis

Trigger: New interview transcript.

Inputs: Transcript plus cumulative customer-insight file.

AI work: Extract needs, buying criteria, language, pain points, objections, supplier perceptions, and new themes; compare with prior interviews.

Human role: Validate important interpretations and decide implications.

Output: Interview summary plus updated theme tracker.

Best autonomy level: 2.

Workflow 9 — Voice-of-Customer Library

Trigger: New source of customer language.

Inputs: Calls, interviews, surveys, reviews, support messages.

AI work: Extract useful verbatim phrases and tag them by theme, stage, and customer type.

Human role: Remove misleading or confidential material.

Output: Searchable customer-language library.

Best autonomy level: 3 for extraction, 1 for external use.

Workflow 10 — Content Idea Generator From Real Customer Signals

Trigger: Weekly content planning.

Inputs: Recent customer questions, objections, search topics, sales notes, market changes.

AI work: Propose content ideas tied to specific customer problems, with audience, angle, evidence, and likely format.

Human role: Choose ideas worth publishing.

Output: Ranked idea list.

Best autonomy level: 1.

Workflow 11 — Content Brief to Draft

Trigger: Approved content brief.

Inputs: Brief, Company Brain, voice examples, source material.

AI work: Draft a structured article, email, script, or post using approved claims and evidence.

Human role: Add experience, opinion, examples, and final voice.

Output: Strong first draft.

Best autonomy level: 1.

Workflow 12 — Content Repurposing

Trigger: Final long-form content approved.

Inputs: Final article, webinar, podcast, or video transcript.

AI work: Produce channel-appropriate short posts, email summary, sales talking points, FAQ ideas, and excerpts without changing the core message.

Human role: Approve what is actually worth publishing.

Output: Repurposing package.

Best autonomy level: 2 or 3.

Workflow 13 — Campaign Builder

Trigger: Campaign objective approved.

Inputs: Audience, offer, objective, positioning, proof, channel constraints.

AI work: Create campaign brief, message hierarchy, content plan, email sequence outline, ad concepts, and measurement plan.

Human role: Approve strategy before asset production.

Output: Campaign blueprint.

Best autonomy level: 1.

Workflow 14 — Prospect Research

Trigger: Qualified prospect meeting scheduled.

Inputs: Company name, participants, CRM history, offer context.

AI work: Research company, recent developments, likely priorities, relevant triggers, and questions to explore.

Human role: Decide hypotheses and call strategy.

Output: One-page prospect brief.

Best autonomy level: 2 or 3.

Workflow 15 — Sales Call Preparation

Trigger: Meeting within defined time window.

Inputs: Prospect brief, opportunity stage, previous notes, offer information.

AI work: Create must-learn questions, hypotheses to test, objections to explore, and relevant proof points.

Human role: Conduct the conversation naturally.

Output: Discovery guide.

Best autonomy level: 2.

Workflow 16 — Post-Call Follow-Up

Trigger: Sales call transcript available.

Inputs: Transcript, opportunity record, offer information.

AI work: Identify stated priorities, objections, commitments, next steps, and draft personalized follow-up.

Human role: Confirm accuracy and send.

Output: Follow-up message plus opportunity update.

Best autonomy level: 2.

Workflow 17 — Proposal First Draft

Trigger: Opportunity reaches proposal stage.

Inputs: Discovery notes, approved scope, pricing rules, case studies, proposal template.

AI work: Draft customer-specific problem statement, recommended approach, scope, timeline, proof, and next steps.

Human role: Own recommendation, price, commitments, and final wording.

Output: Proposal draft.

Best autonomy level: 1 or 2.

Workflow 18 — Customer Onboarding Pack

Trigger: Signed agreement or completed purchase.

Inputs: Customer information, purchased offer, scope, timeline, service policies.

AI work: Prepare welcome message, kickoff agenda, intake summary, requested materials, timeline, and project setup checklist.

Human role: Confirm customer-specific details.

Output: Onboarding package.

Best autonomy level: 2 or 3.

Workflow 19 — Project Status Update

Trigger: Weekly schedule or milestone completion.

Inputs: Recent meetings, action log, milestone status, open risks.

AI work: Summarize completed work, current status, next steps, customer actions, owner actions, and risks.

Human role: Validate anything external-facing.

Output: Internal or customer status report.

Best autonomy level: 2.

Workflow 20 — SOP Builder

Trigger: A recurring process is performed or recorded.

Inputs: Notes, screen recording transcript, example output, known exceptions.

AI work: Convert the process into purpose, prerequisites, steps, checks, exceptions, and outputs.

Human role: Test the procedure against real work.

Output: Standard operating procedure.

Best autonomy level: 1.

Workflow 21 — Customer Support Triage

Trigger: New support request.

Inputs: Customer message, account context, approved knowledge base, escalation policy.

AI work: Classify issue, determine risk level, retrieve relevant information, draft response or escalate.

Human role: Handle exceptions and sensitive cases.

Output: Draft response or escalation brief.

Best autonomy level: 2 initially; higher only for narrow routine categories.

Workflow 22 — Document Summary and Obligation Extraction

Trigger: New contract, proposal, policy, or long document received.

Inputs: Document and review purpose.

AI work: Summarize material terms, dates, responsibilities, financial obligations, risks, and items requiring professional review.

Human role: Verify against source; obtain legal or other expert advice where appropriate.

Output: Structured document brief.

Best autonomy level: 1.

Workflow 23 — Vendor Comparison

Trigger: Two or more proposals available.

Inputs: Vendor proposals and approved evaluation criteria.

AI work: Normalize differences, build side-by-side comparison, identify missing information, and flag material tradeoffs.

Human role: Decide which tradeoffs matter and select vendor.

Output: Comparison matrix plus unanswered questions.

Best autonomy level: 1.

Workflow 24 — Monthly Management Report

Trigger: Month-end data available.

Inputs: Sales, financial, marketing, customer, and operational metrics.

AI work: Identify material changes, trends, anomalies, and items needing explanation; produce draft management commentary.

Human role: Add causal context and decide actions.

Output: Monthly management brief.

Best autonomy level: 2 or 3 for preparation.

Workflow 25 — Quarterly Strategic Review

Trigger: End of quarter.

Inputs: Goals, scorecards, customer insight, competitive intelligence, financial results, major decisions, open risks.

AI work: Summarize what changed, which assumptions held or failed, what worked, what did not, emerging opportunities, and decisions for the next quarter.

Human role: Set direction.

Output: Strategic review memo and next-quarter priorities.

Best autonomy level: 1.

Starter Instructions for Your First AI Roles

These are not “magic prompts.” They are starting instruction sets. Customize them with your Company Brain, tools, authority limits, and actual workflow.

Starter Instruction — AI Chief of Staff

Mission

Help me operate the business with fewer dropped commitments and better preparation. Your job is to organize information, surface priorities, prepare decisions, and reduce administrative friction. You don't make consequential business decisions for me.

Working Principles

1. Distinguish facts, commitments, assumptions, and suggestions.
2. Don't turn a possible next step into an agreed action.
3. When information is missing, say what is missing rather than inventing it.
4. Keep briefings concise enough to use.
5. Surface conflicts between deadlines or commitments.
6. Flag decisions that are waiting on me.
7. When recommending priorities, explain the business reason rather than simply ranking tasks by urgency.

Standard Morning Brief

Prepare:

- Today's three most important outcomes.
- Meetings requiring preparation.
- Commitments due today or overdue.
- Decisions waiting on me.
- Material changes or risks.
- Items that can safely wait.

Keep the brief to one page unless I request detail.

Escalate / Highlight

Always highlight:

- Missed or conflicting commitments.
- Customer or client issues that may affect trust.
- Financial or contractual deadlines.
- Unusual changes in an active project.
- Anything you are uncertain about but believe may be consequential.

Starter Instruction — AI Research Analyst

Mission

Help me become well informed enough to make better business decisions. Favor reliable evidence, current information, and clear uncertainty over polished but unsupported conclusions.

Working Principles

1. Start by restating the research question and scope.

2. Separate verified facts from source claims and your interpretation.
3. Prefer primary sources when they are reasonably available and relevant.
4. Use recent sources when the question is time-sensitive.
5. Label geography, population, and time period for material statistics.
6. Don't create precision the evidence doesn't support.
7. Include disagreement when credible sources materially conflict.
8. End important research with: What We Know / What We Think / What We Still Need to Learn.
9. Provide source information for material claims.

Standard Research Brief

- Research question.
- Executive summary.
- Key findings.
- Supporting evidence.
- Contradictory or limiting evidence.
- Implications for our business.
- Unknowns.
- Suggested next research step, if needed.

Escalate / Highlight

Flag:

- Weak or circular sourcing.
- Important data older than the required period.
- Major claims that appear only in vendor marketing.
- Areas where available evidence cannot support a reliable conclusion.

Starter Instruction — AI Growth Specialist

Mission

Help the company create relevant demand and support sales without replacing human relationships or producing generic high-volume content.

Working Principles

1. Start from approved customer insight, positioning, offers, and proof.
2. Prefer relevance over volume.
3. Don't invent customer pain points, testimonials, case studies, or results.
4. Match the message to the buyer's actual stage and situation.
5. Use customer language where appropriate, but don't mimic private customer statements externally without permission.
6. Avoid hype, empty superlatives, and generic AI-style business language.
7. For campaigns, create or confirm the brief before creating assets.
8. For sales support, distinguish prospect facts from hypotheses.
9. Never imply an offer, discount, guarantee, or contractual term that isn't approved.

Standard Campaign Brief

- Audience.
- Objective.
- Customer problem or trigger.
- Offer.
- Main message.
- Proof.
- Objection.
- Call to action.
- Channels.
- Success measure.

Standard Prospect Brief

- Company snapshot.
- Recent developments.
- Known prior interactions.
- Likely connection to our offer.
- Three hypotheses to test.
- Five questions worth asking.
- One thing we should not assume.

Starter Instruction — AI Operations Coordinator

Mission

Help work move through the company consistently. Convert conversations and decisions into actions, procedures, project visibility, and reliable follow-through.

Working Principles

1. Don't invent owners or due dates when they were not assigned.
2. Distinguish an explicit decision from a discussion point.
3. Keep project status factual.
4. Surface blocked work and overdue commitments.
5. Build procedures from actual practice, not generic best practices, unless asked.
6. When an exception appears, record it rather than forcing it into the normal workflow.
7. Keep documents simple enough to use during real work.

Standard Meeting Output

- Purpose.
- Decisions.
- Commitments.
- Action items.
- Owners.
- Due dates, if stated.
- Open questions.

- Risks.
- Follow-up draft.

Standard SOP Structure

- Purpose.
- Trigger.
- Inputs / prerequisites.
- Steps.
- Quality checks.
- Exceptions.
- Output.
- Last updated.

Starter Instruction — AI Customer Support Specialist

Mission

Help customers receive accurate, useful, and timely responses from approved company information. Protect the relationship by recognizing when a human should take over.

Working Principles

1. Answer from approved knowledge and account information only.
2. Never invent a product capability, policy, delivery date, refund right, or contract term.
3. If the answer is uncertain, say so and escalate.
4. Be concise, respectful, and practical.
5. Don't expose another customer's information.
6. Treat anger, safety concerns, privacy issues, legal threats, unusual financial requests, and policy exceptions as escalation categories unless a specific approved workflow says otherwise.
7. Separate classification from response generation.

Classification Categories

- Routine information.
- Account-specific information.
- Technical or product issue.
- Complaint.
- Billing / refund.
- Policy exception.
- Safety / legal / privacy / security.
- Unknown.

Output

Provide:

- Classification.
- Relevant knowledge used.
- Draft response or escalation recommendation.
- Confidence / missing information.

Starter Instruction — AI Back-Office Analyst

Mission

Reduce the owner's administrative burden by extracting, comparing, summarizing, organizing, and explaining business information. Prepare information for decisions without pretending to hold professional authority you don't have.

Working Principles

1. Preserve source meaning when summarizing.
2. Cite document sections or page references for material findings when possible.
3. Don't silently resolve contradictions between documents.
4. Flag missing information.
5. For financial information, distinguish calculation and preparation from accounting, tax, or investment advice.
6. For contracts, summarize terms and identify review questions rather than declaring a contract "safe."
7. For vendor comparison, use our criteria rather than the vendor's preferred framing.
8. Keep an audit trail for material transformations.

Standard Document Brief

- Document purpose.
- Parties / source.
- Key terms or findings.
- Dates.
- Obligations.
- Financial items.
- Risks or unusual provisions.
- Missing information.
- Items requiring human or professional review.

Practical Prompt Pattern Library

Prompts are not the foundation of the AI department, but good instructions still matter. The patterns below are designed to help you create prompts that behave more like miniature operating procedures.

Pattern 1 — Context + Job + Standard + Boundary

Use this for straightforward recurring work.

Context: Here is the business, customer, project, or source information you need.

Job: Here is the result I need you to produce.

Standard: Here is what a good result must contain.

Boundary: Here is what you must not assume, invent, decide, or do.

Example

Context: This is a transcript from a discovery call with a manufacturing prospect. Use the attached offer description and our sales context file.

Job: Prepare a post-call brief and draft follow-up.

Standard: Capture stated problems, buying criteria, stakeholders, objections, timing, commitments, and unresolved questions. Use the prospect's language where useful.

Boundary: Don't infer budget, authority, urgency, or next steps unless they were stated. Don't add promises or pricing.

This structure is simple enough to remember and strong enough for many business tasks.

Pattern 2 — Extract Before You Interpret

Use this for research, interviews, surveys, meetings, and customer feedback.

Step 1: Extract observations from the source without interpretation.

Step 2: Group related observations into themes.

Step 3: Show evidence supporting each theme.

Step 4: Identify contradictions and outliers.

Step 5: Interpret what the pattern may mean.

Step 6: Clearly label the interpretation as interpretation.

This pattern reduces the chance that AI jumps from one vivid quote to a broad conclusion.

Pattern 3 — Facts / Assumptions / Unknowns

Use this when making decisions.

Ask AI to organize the situation into three columns:

Facts we can verify.

Assumptions we are making.

Unknowns that could change the decision.

Then ask it to propose the fastest reasonable way to reduce the most important unknowns.

This is a useful strategy prompt because it turns uncertainty into work rather than hiding it.

Pattern 4 — Options With Tradeoffs

Use this instead of asking AI to "tell me what to do."

"Create three viable options. For each option, show: expected benefit, cost, implementation effort, risks, reversibility, assumptions, who it works best for, and the information that would cause us to reject it."

You remain the decision maker. AI does the comparative framing.

Pattern 5 — Challenge the Current Plan

Use this when you are worried about confirmation bias.

“Assume our current plan may be wrong. Identify the three assumptions most likely to fail. For each one, show evidence that would weaken it, evidence that would strengthen it, and a low-cost test we could run.”

The purpose is not to be negative. It is to make the plan more robust.

Pattern 6 — Draft, Then Critique Separately

Don't always ask the same pass to create and judge the work.

Pass 1: Create the draft.

Pass 2: Act as a skeptical reviewer. Identify unsupported claims, generic language, missing customer logic, contradictions, and unnecessary length.

Pass 3: Revise only after the critique is accepted.

This separation is especially useful for proposals, reports, and long-form content.

Pattern 7 — Use an Approved Source Set

For customer-facing or factual work:

“Use only the attached approved documents for company facts, product capabilities, pricing, policies, and proof. You may use external sources only for [defined purpose]. If an answer is not in the approved materials, say that it isn't confirmed.”

This pattern is simple and powerful for reducing invention.

Pattern 8 — Create an Escalation Decision

Before asking AI to answer a customer or take an action, ask it to classify the situation.

“First classify this request into one of the following categories. Then state whether the approved policy allows AI handling, requires owner approval, or requires immediate escalation. Don't draft the response until classification is complete.”

Classification before generation is a useful safety pattern.

Pattern 9 — Turn an Example Into a Standard

If you have a piece of work you like:

“Analyze the attached example. Identify the structural choices, tone, level of detail, evidence style, formatting, and decision logic that make it effective. Create a reusable standard from those features. Don't copy the specific content.”

This helps turn your best work into institutional knowledge.

Pattern 10 — Improve the Workflow, Not Just the Output

After using a workflow, ask:

“Review the process we just used. Where did we lack information? Which step created unnecessary rework? Which instruction was ambiguous? What should be added to the Company Brain? Which exception rule should be added before the next run? Recommend no more than three changes.”

This prompt makes the department learn from the work.

Pattern 11 — Shorten Without Flattening

AI often shortens by deleting the details that make writing human.

Use:

“Reduce this by approximately 25 percent. Keep the specific examples, strong opinions, and natural phrasing. Remove repetition, throat-clearing, generic transitions, and sentences that merely restate the prior sentence. Don't make the tone more

corporate.”

This is useful for maintaining voice while editing efficiently.

Pattern 12 — Humanize Without Adding Fake Personality

Use:

“Revise this so it sounds like an experienced person explaining the subject to another capable person. Use natural sentence-length variation, concrete examples, and direct transitions. Remove inflated claims, generic motivational language, unnecessary headings, and stock phrases. Don't add personal anecdotes that were not provided.”

The last sentence matters. Humanized language should not require invented life experience.

Pattern 13 — Build a Reusable Template From a Finished Deliverable

“Study this finished deliverable and separate the reusable structure from the project-specific content. Create a template with instructions under each section explaining what belongs there, what evidence is required, and what mistakes to avoid.”

This is how one good project becomes a repeatable operating asset.

Pattern 14 — Compare Versions Precisely

“Compare Version A and Version B. Create a table showing substantive changes in scope, price, timing, obligations, assumptions, and risk. Ignore cosmetic wording changes unless they alter meaning. Quote or point to the relevant source location for each material difference.”

This is useful for proposals, agreements, policies, and client documents.

Pattern 15 — Prepare Me, Don't Script Me

For meetings:

“Prepare me for this conversation. Give me the objective, what I should know, three hypotheses to test, five questions worth asking, signals to listen for, and one assumption I should avoid. Don't write a full script.”

This preserves presence and judgment.

Common Failure Modes and How to Fix Them

Building an AI department is less like installing a machine and more like improving a process. The first version will expose gaps. That is useful. The following problems are common enough that it helps to recognize them early.

Failure Mode 1 — The Output Is Polished but Generic

What It Looks Like

The writing is grammatically clean, logically organized, and almost completely interchangeable with what a competitor might say.

You see phrases such as “drive meaningful growth,” “unlock new opportunities,” “streamline operations,” or “deliver exceptional value.” Nothing is technically wrong. Nothing sounds like your company either.

Why It Happens

The system doesn't have enough specific context, examples, customer language, or point of view.

It is filling the empty spaces with statistically safe business language.

Fix It

Don't respond by adding, “Make it more compelling.”

Add better raw material.

Give the model:

- Specific customer phrases.
- Real examples.
- Clear opinions.
- Approved proof.
- Strong samples of your writing.
- Phrases you refuse to use.

Ask it to remove any sentence that could be used unchanged by five competitors.

The general lesson: generic output is often a context problem before it is a writing problem.

Failure Mode 2 — You Are Spending More Time Managing AI Than You Saved

What It Looks Like

The workflow is clever but fragile. You have to check three systems, move files manually, restart failed automations, and remember special commands.

The “automation” has become another job.

Why It Happens

You optimized for technical possibility instead of operating simplicity.

Fix It

Strip the workflow back to the business outcome.

Ask:

- Which step creates real value?
- Which steps exist only because two tools don't work well together?
- Could a manual handoff be faster than maintaining this integration?
- Could one general tool replace two specialized tools?

A one-person business should prefer a workflow that saves 80 percent of the time and is easy to maintain over one that saves 95 percent and breaks every other Thursday.

The general lesson: maintenance cost is part of automation cost.

Failure Mode 3 — The AI Keeps Forgetting How You Work

What It Looks Like

One day the output is excellent. The next day it ignores your format, uses outdated information, or makes a recommendation inconsistent with your strategy.

Why It Happens

Important context lives in conversation history rather than in maintained source documents.

Fix It

Move durable knowledge into the Company Brain.

Create explicit files for:

- Current strategy.
- Offers.
- Brand rules.
- Customer definitions.
- Workflow standards.
- Approved proof.

Then tell the AI which source is authoritative.

Don't make a chat thread responsible for being your company's memory forever.

The general lesson: important context should be stored intentionally, not remembered accidentally.

Failure Mode 4 — Every Workflow Has Become a Giant Prompt

What It Looks Like

Your instructions are several pages long. Every exception creates another paragraph. Nobody — including you — wants to edit the thing.

Why It Happens

You are trying to solve process design with instruction volume.

Fix It

Break the workflow into stages.

Instead of one giant prompt for “research, analyze, recommend, write, verify, format, and send,” create smaller steps with clear handoffs.

Put stable facts in context documents. Put recurring process rules in the workflow. Put output requirements in templates. Put exceptions in a short escalation list.

The general lesson: if one prompt is carrying the entire operating system, the operating system is underdesigned.

Failure Mode 5 — You Automated a Bad Process

What It Looks Like

The old process was slow. The new process is fast. The output still doesn't matter.

You now produce an unread report automatically, generate content nobody needs more efficiently, or follow up on leads that were never worth pursuing.

Why It Happens

Automation began before the business purpose was questioned.

Fix It

Return to the activity and ask:

- Who uses this output?
- What decision does it support?
- What happens if we stop doing it?
- Could the desired outcome be achieved more simply?

If the work fails those questions, eliminate it.

The general lesson: automation should come after usefulness.

Failure Mode 6 — The System Is Too Confident When It Is Missing Information

What It Looks Like

AI fills in unstated dates, assumes the customer's budget, guesses a product capability, or creates a plausible explanation for a number.

Why It Happens

Most AI systems are optimized to be helpful. If the instruction doesn't reward uncertainty, the model may resolve ambiguity instead of exposing it.

Fix It

Create explicit missing-information behavior.

For example:

"If any required input is missing, list it under Missing Information. Don't infer pricing, dates, customer commitments, product compatibility, legal meaning, or factual claims."

In structured workflows, add a preflight check that confirms inputs before drafting begins.

The general lesson: teach the system that stopping is sometimes the correct result.

Failure Mode 7 — You Do Not Know Whether the Workflow Is Actually Better

What It Looks Like

The workflow feels modern. You like having it. But you cannot say whether it saves time, improves quality, or supports revenue.

Why It Happens

There was no baseline.

Fix It

Before improving a process, capture a rough "before" measure.

How many minutes does it take?

How long does the customer wait?

How often is rework required?

How often do you miss a step?

Then measure again after several runs.

You don't need perfect analytics. You need enough evidence to distinguish leverage from novelty.

The general lesson: what feels efficient and what is efficient are not always the same.

Failure Mode 8 — You Have Too Many AI Roles

What It Looks Like

You have twelve specialized assistants, six projects, nine agents, and no idea which one should handle a routine task.

Why It Happens

Role design became entertainment.

Fix It

Merge roles until each one has recurring work.

Most one-person businesses can go a long way with three to six roles. Specialization should solve a context or process problem, not satisfy your desire for a fancy org chart.

If two roles use nearly the same context, produce similar outputs, and follow similar rules, they may be one role.

The general lesson: organizational complexity is still complexity, even when the employees are digital.

Failure Mode 9 — AI Has Flattened Your Voice

What It Looks Like

Everything is clean. Everything is sensible. Nothing sounds quite like you anymore.

Why It Happens

AI is drafting too close to the final version, or the editing process rewards smoothness more than personality.

Fix It

Change the division of labor.

You might provide the opening, argument, examples, or rough voice notes yourself and use AI for structure and tightening.

You might ask AI for an editorial memo instead of a rewrite.

You might keep a list of phrases that sound fake coming from you.

Most important, preserve specific details. Human writing often feels human because it contains choices, not because it contains imperfections.

The general lesson: don't outsource the parts of communication that are the reason people listen to you.

Failure Mode 10 — You Are Letting AI Decide Before It Has Earned the Right

What It Looks Like

The system recommends a vendor, rejects a lead, prioritizes customers, or chooses a strategic direction and you begin accepting the answer because it arrives neatly formatted.

Why It Happens

Preparation and decision authority have blurred together.

Fix It

Move the workflow down the autonomy ladder.

Ask the AI to prepare options, evidence, tradeoffs, and unknowns.

Reserve the choice for yourself until you have strong reason to do otherwise.

For high-consequence decisions, keep human ownership permanently.

The general lesson: a well-prepared decision is not the same thing as a delegated decision.

Failure Mode 11 — Your Automation Is Creating Customer Distance

What It Looks Like

Response time is excellent, but customers receive messages that feel oddly impersonal. Important customers start replying with “Can I just talk to you?”

Why It Happens

You optimized for speed and consistency without defining relationship moments.

Fix It

Map the customer journey and mark the moments that deserve you personally.

Examples may include:

- First meaningful sales conversation.
- Major proposal discussion.
- Kickoff.
- Serious complaint.
- Renewal.
- Strategic review.
- A customer’s significant success or setback.

Use AI to prepare for those moments, not to remove you from them.

The general lesson: the purpose of automation is to free relationship time, not eliminate relationships.

Failure Mode 12 — The Company Brain Has Become a Junk Drawer

What It Looks Like

The knowledge base contains duplicate files, outdated versions, old offers, draft policies, and documents you no longer trust.

Why It Happens

Information was added but never governed.

Fix It

Use three simple status labels:

Current — approved for active use.

Reference — useful historical material.

Draft — not authoritative.

Once a quarter, remove or archive material that is no longer needed.

For critical topics, identify one source of truth.

The general lesson: retrieval quality is limited by knowledge quality.

Failure Mode 13 — You Are Automating Before You Understand the Work

What It Looks Like

The workflow repeatedly encounters exceptions you did not anticipate. You keep adding rules. The process never stabilizes.

Why It Happens

You tried to automate a process you had not observed carefully enough.

Fix It

Run the workflow manually with AI assistance for ten or twenty cases.

Record:

- Common inputs.

- Common outputs.
- Edge cases.
- Missing information.
- Judgment points.
- Rework.

Only then decide which parts deserve automation.

The general lesson: manual repetition is often the research phase of automation.

Failure Mode 14 — You Have No Exit Plan When a Tool Changes

What It Looks Like

A vendor changes pricing, removes a feature, or alters an integration and an important business process stops.

Why It Happens

The workflow exists only inside the vendor's interface.

Fix It

Document the process in vendor-neutral language.

Keep your important source files in formats you can export. Know which inputs and outputs matter. Avoid putting irreplaceable knowledge only inside a proprietary workflow.

You don't need to be ready to switch platforms tomorrow. You should be able to explain how the business process works without pointing at a screen.

The general lesson: own the workflow even when you rent the technology.

Failure Mode 15 — You Keep Adding Instead of Improving

What It Looks Like

Every month brings another workflow. The first ones still have obvious problems, but building something new is more interesting than fixing them.

Why It Happens

Creation feels like progress. Maintenance feels ordinary.

Fix It

Use the weekly AI Operations Review.

For every new workflow you add, improve or retire one existing workflow.

Measure the department by reliability and business impact, not by the number of automations.

A one-person company doesn't need a hundred AI workflows. It needs a small set of workflows that quietly make the company better.

The general lesson: depth beats breadth after the experiment phase.

The One-Page AI Department Operating Manual

Once your system is running, compress the essentials onto one page. This is the document you should be able to look at after a busy month and immediately remember how the department is supposed to work.

Our AI Department Exists To

Complete this sentence:

“Our AI department exists to _____ so that I can spend more of my time on _____.”

A strong answer names both the work you want AI to absorb and the human work you are trying to protect.

Our Current AI Roles

Role 1:

Primary job:

Role 2:

Primary job:

Role 3:

Primary job:

Optional additional roles:

Our Three Most Valuable Workflows

Workflow 1:

Trigger:

Output:

Human checkpoint:

Workflow 2:

Trigger:

Output:

Human checkpoint:

Workflow 3:

Trigger:

Output:

Human checkpoint:

Our Source of Truth

Core Business Context location:

Approved product / offer information:

Approved customer information:

Approved brand guidance:

Approved policies:

Current project context:

Our Five Non-Negotiable Rules

1.

- 2.
- 3.
- 4.
- 5.

What AI May Do Without Approval

-
-
-

What Always Requires Approval

-
-
-

What Must Be Escalated Immediately

-
-
-

Weekly Review

Day / time:

Scorecard location:

Questions:

- What saved real time?
- What created rework?
- What failed?
- What context changed?
- What should we simplify?
- What is ready for greater autonomy?

Current 30-Day Improvement Goal

"This month we are improving _____."

If you cannot fit the department's operating logic on one page, simplify until you can.

Final Reader Checklist: Is This a Department Yet?

Use this checklist after thirty to ninety days. You don't need every box checked. The point is to see whether your AI use has moved from occasional assistance toward a real operating capability.

Strategy

- I can explain why we use AI in one sentence.
- I know which human work I am trying to protect or expand.
- I am not using automation simply because it is possible.

Work Audit

- I have identified my most repetitive work.
- I know which work should remain human-owned.
- I have eliminated at least one low-value activity instead of automating it.

Roles

- My recurring AI work has clear functional homes.
- Each important AI role has a mission and authority boundary.
- I know when each role must escalate to me.

Company Brain

- AI has access to current, approved business context.
- Customers, offers, positioning, and policies are documented.
- Important source documents have a clear current version.
- Sensitive information is handled deliberately.

Workflows

- At least three recurring activities have defined workflows.
- Each workflow has a trigger, inputs, process, human checkpoint, output, and destination.
- I have defined what “done” looks like.
- Important exceptions are documented.

Automation

- Every automated workflow has an appropriate autonomy level.
- I know what the system may do without asking me.
- I can stop the automation quickly.
- I have enough logging to understand what happened if something goes wrong.

Quality

- Important facts, numbers, dates, and commitments are verified.
- High-risk work receives higher human oversight.
- I track common errors or failure modes.
- AI-generated customer communication still sounds like the business.

Performance

- I have a rough baseline for important workflows.
- I can name at least one workflow that saves meaningful owner time.

■ I can name at least one workflow that improves quality or consistency.

■ I review the department at least monthly.

The Most Important Test

■ I spend more time on work that requires my judgment, relationships, taste, creativity, or leadership than I did before building the system.

If that final box is not becoming more true over time, the department needs redesign.

That is the point of the whole exercise.

90-Day AI Department Maturity Roadmap

The first thirty days give you a department. The next sixty make it dependable.

Don't spend Months 2 and 3 chasing the latest tools. Use the time to deepen what is already working.

Days 31–45 — Stabilize

Review your three core workflows and ask what still requires unnecessary effort.

For each workflow:

- Remove redundant steps.
- Improve weak inputs.
- Update the relevant Company Brain files.
- Clarify the output template.
- Document the most common exceptions.
- Measure owner time again.

At this stage, avoid adding autonomy simply to make the workflow feel advanced. Reliability comes first.

A good question is: Would I trust this process on a week when I am unusually busy?

If the answer is no, find the part that still depends on remembering a special step.

Days 46–60 — Connect

Now look for handoffs between workflows.

Can the meeting-to-action workflow update the weekly CEO review automatically?

Can customer interview analysis feed the voice-of-customer file used by the Growth Specialist?

Can the prospect research workflow use the same account context later used for proposal preparation?

Can the support system send recurring product questions into the monthly customer insight brief?

The purpose of connection is not technical elegance. It is to stop recreating information.

One useful rule is: important information should be captured once and reused many times.

Be careful about creating a chain where an early error silently spreads into five downstream outputs. Build checkpoints around information that becomes authoritative.

Days 61–75 — Graduate One Workflow

Choose one low- or moderate-risk workflow that has been consistently reliable.

Review its track record.

Ask:

- How often do I change the output materially?
- What errors still occur?
- Are the exceptions recognizable?
- Can an error be reversed?
- Does the workflow know when required information is missing?

If the evidence is strong, move the workflow up one level on the Autonomy Ladder.

Don't jump from "AI prepares" to "AI acts freely."

A sensible graduation may be as small as moving from manual start to scheduled start, or from owner approval of every internal report to exception-only review.

Record the change and watch it closely for several weeks.

Days 76–90 — Add One New Capability

Only after the core is stable should you expand.

Return to your AI Opportunity Audit and choose one additional capability that addresses a real bottleneck.

You might add:

- Customer support triage.
- A monthly management report.
- Proposal drafting.
- Vendor comparison.
- Content repurposing.
- Market monitoring.

Build the new capability using the same discipline as the first three: role, context, workflow, authority, quality check, measurement.

Don't assume you have “learned AI” and can now skip the design steps. The process is what makes the new capability trustworthy.

The Day-90 Review

At the end of ninety days, take one hour and answer these questions:

1. Which workflow has created the most real value?
2. Which workflow looked promising but has not earned its place?
3. What work do I now do less often or not at all?
4. What customer experience has improved?
5. What information do I understand better than I did three months ago?
6. What errors or risks have appeared?
7. Which role has become most useful?
8. Where am I still the unnecessary bottleneck?
9. Where is my judgment essential and worth protecting?
10. What should the AI department become over the next quarter?

Then write one paragraph describing how the business operates differently than it did ninety days earlier.

If your answer is mostly about software — “I now use five AI tools” — keep working.

If your answer is about business behavior — “every sales call is prepared, every meeting produces actions, customer feedback gets analyzed monthly, proposals take half the time, and I spend Friday mornings on strategy instead of administration” — you have built something meaningful.

That is the standard.

Reader License

Copyright © 2026 James Inglis. The purchaser or reader may copy and adapt these tools for personal use or for use inside a business they own or manage. The toolkit may not be resold, republished, or redistributed as a standalone product.