

PORT JACKSON WORKSHOP · FREE COURSE · DO IT YOURSELF SERIES

Prompts Are Not Agents

How to build a personal AI operating system that actually runs — in three layers, with approval steps — instead of collecting prompt lists that promise 24/7 and deliver once

WHO THIS IS FOR

Solo operators and small-business owners who already pay for an AI assistant, have saved prompt lists that disappointed them, and want workflows that run on a schedule with their decisions kept in their hands.

WHAT YOU WILL HAVE AT THE END

A sorted inventory of repeating work; a standing-context page; a correction log; three chat workflows run daily for a week; two scheduled workflows delivering unattended; one automation job with an approval tap; and a monthly review with hours saved measured, not guessed.

Christine de Vries · Port Jackson Group · v1.0 · 3 September 2026 · Free to read, copy and share unchanged with this notice. This course names no company, coach or product. Every figure is an estimate to check against your own numbers. It is general information, not financial, legal or tax advice. Nothing is sold inside it.

How to use this course

Five modules; the first two are paper, the last three are building. Layer A needs no paid tool. Layer B needs the assistant's paid tier (scheduled tasks and connectors). Layer C needs an automation platform (a free self-hosted option exists). Every lesson ends with a 'Do this now' box. Prices are September 2026 estimates.

Course map

Module	Lessons	You finish with
Inventory	What repeats, and what it costs · Cheap mistakes and expensive ones · The three layers	A four-column inventory with every task tagged by layer and by mistake cost.
Context	The standing-context page · Connecting sources deliberately · The correction log	A standing-context page loaded into a project; a connection register; a correction log with two entries.
Layer A — run it by hand first	Morning brief and goal plan · Weekly plan and Friday review · Design before you automate	Three workflows run for five days each with corrections logged; a six-field design for three more.
Layer B — scheduled inside the assistant	The research digest · Mail triage with approval · Alert rules	A weekly research digest and a twice-daily mail triage running on schedule, drafts only, alerts written.
Layer C — automation with a tap	The first job · The approval step · The monthly review	One job running with an approval step; a first monthly review with hours saved measured.

Module 1 — Inventory

Automation fails when it is aimed at the wrong work. One page decides what is worth automating and how carefully.

Lesson 1.1 — What repeats, and what it costs

List every task you do more than once a week: reading mail, planning the day, chasing invoices, checking sales, researching a topic, writing the same three replies. Four columns: task, the tool it lives in, minutes per week, and what a mistake in it would cost you.

DO THIS NOW

Fill W1. Aim for 15–30 rows. If a task is not on the page, it does not get automated.

Worksheet: W1 Inventory (Appendix)

Lesson 1.2 — Cheap mistakes and expensive ones

A wrongly filed newsletter costs nothing; a wrongly sent quote costs a client. Tag each row cheap or expensive. Cheap-mistake tasks can run with approval after the fact; expensive ones get approval before anything leaves your account. This tag decides where the human sits in every workflow you build.

DO THIS NOW

Tag every row on W1.

Worksheet: W1 Inventory (Appendix)

Lesson 1.3 — The three layers

Layer A — a saved prompt in a project with standing context; runs when you open it; free. Layer B — the same prompt scheduled inside the assistant with results delivered to you; needs the paid tier. Layer C — a job in an automation platform that watches a source, calls the model, writes somewhere and waits for your tap; needs the platform and connectors.

The word 'agent' belongs only to Layers B and C. A prompt that runs once when pasted is a prompt, however it is titled.

DO THIS NOW

Write A, B or C beside every row on W1. Nothing goes to C until it has run in A for two weeks.

Worksheet: W1 Inventory (Appendix)

MODULE 1 — YOU ARE DONE WHEN

A four-column inventory with every task tagged by layer and by mistake cost.

Module 2 — Context

Prompt lists make you retype your goals eight times. Write them once, in one place, and every workflow reads them.

Lesson 2.1 — The standing-context page

One document: who you are and what you do; goals for the quarter; priorities in order; the important-people list and what each is to you; the tools you use; your writing voice; and the approval rules — 'draft replies to anyone; never send; never agree a price or a date.' Load it into a project so every conversation starts with it.

DO THIS NOW

Write W2 and load it. Test by asking the assistant 'what are my top three priorities?' — it should answer from the page.

Worksheet: W2 Standing context (Appendix)

Lesson 2.2 — Connecting sources deliberately

The assistant can only read what is connected. Connect the business inbox, not a personal one; the calendars and folders the workflows need, not everything. Read the data-use terms once. Write on W3 what is connected and why — this is the answer to 'what does the AI see?' when anyone asks.

DO THIS NOW

Connect one inbox. Fill W3.

Check: Could you explain to a client, in one sentence, what the AI can read? If not, connect less.

Worksheet: W3 Connection register (Appendix)

Lesson 2.3 — The correction log

Every time a workflow gets something wrong, one line: date, workflow, what was wrong, what context would have prevented it. After two weeks the log is your to-do list for the standing-context page. This is the step every prompt list skips, and it is the difference between a system that improves and a prompt that disappoints.

DO THIS NOW

Start W4 now. Make two entries this week, even small ones.

Worksheet: W4 Correction log (Appendix)

MODULE 2 — YOU ARE DONE WHEN

A standing-context page loaded into a project; a connection register; a correction log with two entries.

Module 3 — Layer A — run it by hand first

Three workflows in a chat, daily, for a week. This is where you learn what the AI gets wrong before you let it run alone.

Lesson 3.1 — Morning brief and goal plan

Morning brief prompt (paste each morning with your calendar and any mail summary): 'Using my standing context, give me a five-line brief: the three things that matter today, deadlines in the next seven days, and any drafts worth preparing. Separate what you can prepare from what needs my decision.'
Goal plan prompt (once per goal): 'Turn this goal into milestones with dates, recurring actions, dependencies and one measurable outcome each. Tell me which recurring actions could become scheduled workflows.'

DO THIS NOW

Run the brief five mornings running. Log every correction on W4.

Worksheet: W4 Correction log (Appendix)

Lesson 3.2 — Weekly plan and Friday review

Keep the plan in a document, not in the chat. Sunday: 'Build my week from these projects, deadlines and available hours; flag conflicts, overdue items and anything to batch.' Friday: paste the plan back with what was done: 'Compare planned versus completed; give me three improvements for next week.' The comparison only works because the plan persisted — a chat forgets; a document does not.

DO THIS NOW

Do one full Sunday-to-Friday cycle.

Worksheet: W4 Correction log (Appendix)

Lesson 3.3 — Design before you automate

Before any scheduling, specify each workflow in six fields: trigger, required context, expected output, verification method, alert condition, human approval point. Prompt: 'Using my standing context and this inventory, design three workflows with those six fields each.' The design is the deliverable; the running system comes later, and only for designs that survived a week of Layer A.

DO THIS NOW

Complete W5 for three workflows.

Worksheet: W5 Six-field spec (Appendix)

MODULE 3 — YOU ARE DONE WHEN

Three workflows run for five days each with corrections logged; a six-field design for three more.

Module 4 — Layer B — scheduled inside the assistant

Two workflows that deliver without you opening anything. This is the first thing that deserves the word 'agent'.

Lesson 4.1 — The research digest

Schedule for Monday 6am: 'Research these topics [list] using web search. Compare against last week's digest [attached document]. Report what changed, mark each item verified (source seen) or reported (single source), and flag any competitor, price or regulation change.' Keep each week's digest as a document so 'changed' has a baseline. Without the baseline the workflow cannot flag anything — which is why the prompt-list version never worked.

DO THIS NOW

Schedule it. After the first run, save the output as the baseline document and attach it to the task.

Worksheet: W5 Six-field spec (Appendix)

Lesson 4.2 — Mail triage with approval

Schedule 7am and 1pm: 'Review mail since the last run. Classify urgent / this week / can wait using my important-people list. Draft replies for urgent items. Never send. Alert me if a listed person, a payment word or a legal word appears.' The first correction most people log is a supplier invoice filed as 'can wait'; add the rule and move on.

DO THIS NOW

Run for one week. Compare each morning against the two things you already knew were urgent.

Check: Did anything get sent? It must not. If it did, the approval rule is missing from the context.

Worksheet: W4 Correction log (Appendix)

Lesson 4.3 — Alert rules

An alert is a condition, not a mood. Write five: a named person writes; a payment or invoice word appears; a metric moves more than X%; a deadline is within 48 hours; a competitor is named in the digest. Test one by triggering it yourself.

DO THIS NOW

Write five alert rules on W5 and test one.

Worksheet: W5 Six-field spec (Appendix)

MODULE 4 — YOU ARE DONE WHEN

A weekly research digest and a twice-daily mail triage running on schedule, drafts only, alerts written.

Module 5 — Layer C — automation with a tap

One job in an automation platform that watches a source, calls the model, writes one line somewhere, and waits for you before anything leaves. Then the monthly review that keeps the system honest.

Lesson 5.1 — The first job

Platform: a self-hosted open-source workflow tool on a small server (~A\$8 a month) or a cloud plan from ~A\$30. Job: on a schedule, read a source (a sales export, a spreadsheet, a form), send it to the model with a short prompt and your thresholds, and post one line to a channel or a sheet. That is the whole job. Add nothing else until it has run for a week.

DO THIS NOW

Build it using W6. Watch it run five times.

Worksheet: W6 Job blueprint (Appendix)

Lesson 5.2 — The approval step

Any action that reaches a third party — send, pay, book, delete — gets a 'wait for approval' step: the job prepares the action, sends you the preview, and does nothing until you tap. In year one there are no exceptions, including for cheap-mistake tasks, because the cost of learning where the line is exceeds the minutes saved.

DO THIS NOW

Add the approval step to the job before it is allowed to write anywhere but a sheet.

Worksheet: W6 Job blueprint (Appendix)

Lesson 5.3 — The monthly review

Once a month: what ran, how many corrections each workflow needed, hours saved (measured against the inventory's minutes column), hours spent running it. Promote workflows that needed no corrections; demote or retire any corrected more than four times. Publish the honest net number to yourself. Most people find four hours a week net, not the forty a headline promised — and four hours a week is 200 a year.

DO THIS NOW

Complete W7 at the end of month one. Retire at least one thing.

Worksheet: W7 Monthly review (Appendix)

MODULE 5 — YOU ARE DONE WHEN

One job running with an approval step; a first monthly review with hours saved measured.

Appendix — Worksheets

W1 Inventory

Task	Tool	Min/week	Mistake cost (cheap/expensive)	Layer A/B/C

W2 Standing context

Section	Content
Who I am / what I do	
Goals this quarter	
Priorities in order	
Important people and what each is to me	
Tools	
Writing voice	
Approval rules	Draft anything; never send; never agree a price or a date

W3 Connection register

Source	Connected (date)	Read-only?	Why	Workflows using it

W4 Correction log

Date	Workflow	What was wrong	Context that would have prevented it

W5 Six-field spec (one per workflow)

Field	Workflow 1	Workflow 2	Workflow 3
Trigger			
Required context			
Expected output			
Verification			
Alert condition			
Human approval point			
Layer / cost			

W6 Job blueprint

Step	Setting
Schedule	
Source read	
Prompt and thresholds	
Output written to (sheet/channel)	
Approval step before any external action	required
Failure alert	

W7 Monthly review

Workflow	Ran (times)	Corrections	Hours saved	Hours spent	Promote / keep / demote / retire
Net hours this month					

Honest limits

- Feature names (projects, scheduled tasks, connectors) and tier prices change; the layer logic does not. Check your assistant's current plan page.
- Connecting a mailbox or drive to any AI service is a data decision; read the terms and connect the minimum.
- No workflow in this course sends, pays, books or deletes without a human tap. Keep it that way for a year.
- Hours saved are real only after a month of measuring; publish nothing before then, to yourself or anyone else.