

rpi-app-workbook

DEPLOYED

Paper-first personal reflection workbook with daily auto-print and cloud sync

auto print/printing needs to be left private
Purpose of this is to have center for human in the loop function

Stack: FastAPI + Postgres + Typst + Ollama (qwen3.5:4b) + CUPS printer + Cloudflare Tunnel

Live: <https://holm.team>

WHAT IT DOES

needs to be public facing, private functions such as this workbook needs to be post login. maybe a store front?
no longer philosophy, now just projects
what is CUPS server for printing?
help make

Transforms public-domain philosophy texts and scanned books into printable workbooks through a database-first composition system. Users fill in pages by hand, then scan them back for OCR indexing. Auto-fires Daily Brief (2-page workbook + schedule) every day at 04:00 PT via K8s CronJob with auto-duplex printing to CUPS printer.

CURRENT STATE

v0.6 architecture live; 1,114 Python files; 117 Project Gutenberg works indexed via scraper (uses Ollama qwen3.5:4b for theme classification); 46 curated theme questions; Postgres deployed at workbook-db.workbook.svc.cluster.local:5432; Typst v0.14.2 rendering; Cloudflared tunnel ID d6b19c36-56ba-4268-a9d8-38e5729f0674; auth via bcrypt sessions; image build workaround uses crane push to t1.sh (24h ephemeral) due to Forgejo 60s timeout on >100MB images

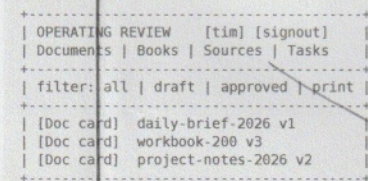
KEY PATHS

- /Users/tim/Documents/HoIm/rpi-app-workbook/webapp/main.py
- /Users/tim/Documents/HoIm/rpi-app-workbook/generator/compose_workbook.py
- /Users/tim/Documents/HoIm/rpi-app-workbook/templates/workbook_v4.typ
- /Users/tim/Documents/HoIm/rpi-app-workbook/k8s/daily-brief-cronjob.yaml

more uses but not sure

what is a format companies use to help executives make decisions?
each paper is for a very good purpose for drafting and creating?

WEBPAGE PREVIEW



Wasted space is being displayed to manipulate the exec?

remove this section

DECISIONS, PENDING, OPEN QUESTIONS

Key decisions

- Database is source of truth—PDF render is pure SQL query, no LLM at render time
- Daily Brief composition automatic at 04:00 PT; prints to duplex queue without approval
- Scraper ingests themes via Ollama (not user input), embeds via nomic-embed-text
- Cloudflare tunnel instead of direct K8s service for SSL termination
- Page-as-document architecture (v0.6): QR codes link to per-page routes for annotation

Pending

- Per-page webapp routes (/page/{serial}) so QR codes work end-to-end
- Stacked-pages homepage (200 pages as list with Preview + upload + comment per page)
- Books reader at /books/{work_id} with scrollable view and #passage-N anchors
- Front matter (cover, foreword, TOC) in v4 template; currently only renders 4-card pages
- Move workbook pod off rpi-8 (single ARM core) to cm3588-plus for better speed
- Fix auth.ensure_schema() ordering bug (counts users before creating table)

Open questions

- Should QR code footer be image or text URL?
- How to handle PDF reflow when user adds handwritten annotations during scan? add as comments to side of document then ollama
- Metric for measuring 'reflection quality' from scanned pages?

38 total pages

276

image -> scan phone take pic upload

needs testing and use of an local text model. ocr model.

API's !!!

MARGIN NOTES (WRITE HERE, IN MARGINS, OR ANYWHERE ELSE)