

AI WORKFLOW PRODUCTS

Workflow Blueprint

Manual Weekly Reporting — an illustrative sample of the deliverable you receive.

Deliverable: single PDF, 10–16 pages
Turnaround: 72 hours from a complete brief
Scope: one workflow · one team · up to five tools · no calls
Independent AI systems practice.

Illustrative example. Based on a representative agency reporting workflow. Company details, volumes, time estimates and expected savings are hypothetical assumptions provided to demonstrate the format. This is not a real client case and does not represent measured customer outcomes. Every quantitative figure below is an illustrative assumption.

Illustrative example. Based on a representative agency reporting workflow. Company details, volumes, time estimates and expected savings are hypothetical assumptions provided to demonstrate the format. This is not a real client case and does not represent measured customer outcomes. Every quantitative figure below is an illustrative assumption.

SECTION 1

Executive summary

A hypothetical 12-person marketing agency assembles a weekly client performance report by hand. A coordinator collects attachments from email, consolidates figures into a spreadsheet, validates the numbers, writes a summary, has it reviewed, and distributes it. In this illustrative scenario the process is assumed to take roughly 5–6 hours every week and to be error-prone at the consolidation and validation steps.


This blueprint proposes a mostly-automated pipeline with two human approval points. Under the stated assumptions, hands-on time could fall to an estimated ~1 hour per week while consistency improves. **All figures are illustrative assumptions, not measured results.** Confidence in the approach is medium-high; confidence in the effort estimate is medium until data formats are confirmed.

Assumed metric (illustrative)	Today	Future state
Hands-on time per week	~5–6 hours	~1 hour
Error-prone manual steps	3 of 6	0 (validated + approved)
Human approval points	1 (ad hoc)	2 (defined)

SECTION 2

Current-state map

How the workflow is assumed to run today — six sequential steps, three of them manual and high-effort.



SECTION 3


Bottleneck analysis

Step	Bottleneck	Why it costs
Attachment collection	Manual search across email threads	Missed or duplicated files; slow start
Consolidation	Copy-paste between formats	Biggest time sink and error source
Validation	No defined rule set	Errors caught late and inconsistently
Review	Waits on one reviewer	A queueing delay, not a work delay

SECTION 4

Future-state map

A mostly-automated pipeline. The automated path is shown in accent; two human approval gates are outlined distinctly. Numbers are never generated by AI.



SECTION 5

Recommended tools & integrations

Keep the existing stack — no rip-and-replace. Up to five tools, in line with Blueprint scope.

Capability	Illustrative option	Notes
Attachment retrieval	Email API or no-code connector	Rules by sender / subject / label
Structured store	Spreadsheet or lightweight database	Single source of truth for figures
Validation	Formulas or a small script	Range, total, and missing-field checks
Summary draft	An LLM step	Receives validated aggregates only — no PII
Distribution	Existing email / chat	Automated once approved

SECTION 6

AI vs deterministic steps

Deterministic (reliable, repeatable)	AI-assisted (always human-approved)
Retrieval, ingest, validation rules, distribution — and all figures.	The narrative summary draft only. Reviewed before it is sent. AI never generates numbers.

SECTION 7

Human approval points

Gate	What is checked	Blocks until
1 — Data sign-off	Validated figures are correct and complete	Resolved before drafting
2 — Summary sign-off	Narrative matches the data and reads well	Approved before distribution

SECTION 8

Errors & controls

- Missing-attachment alert if an expected sender does not report in time.
- Validation flags out-of-range values and mismatched totals, and blocks approval gate 1 until resolved.
- Every run is logged for traceability.
- A manual fallback path runs the old process if automated retrieval fails.

SECTION 9

Security & data risks

In this illustrative design, client data stays within the existing tools. No new third party holds raw data except the LLM step, which receives only validated, non-sensitive aggregate figures — no personal data. Client identifiers are redacted in the summary-draft prompt. Access is limited to the coordinator and reviewer.

SECTION 10

Implementation sequence

Step	Build activity
1	Define retrieval rules and the validation rule set
2	Stand up the structured store and ingest
3	Add validation and approval gate 1
4	Add the AI summary draft and approval gate 2
5	Automate distribution
6	Run in parallel with the manual process for two cycles, then cut over

SECTION 11

30-day roadmap

Week	Focus
Week 1	Retrieval rules + structured store
Week 2	Ingest + validation + approval gate 1
Week 3	AI summary draft + approval gate 2 + distribution
Week 4	Parallel run, tune, and cut over

SECTION 12

Estimated effort & potential time savings

Illustrative estimates only. Build effort: ~2–4 focused days (no-code-leaning) up to ~1 week (scripted). Ongoing: ~1 hour/week versus ~5–6 hours today, an assumed ~4–5 hours/week recovered. Actual figures depend on tool access and data cleanliness and are not guaranteed.

SECTION 13

Assumptions & confidence

Assumes attachments arrive in consistent formats from a stable set of senders and that the team keeps its existing stack. Confidence: medium-high on the approach, medium on effort until formats are confirmed. Where confidence is lower, this blueprint flags it rather than overstating.

This is a sample. A real Workflow Blueprint is written from your brief, scoped to your named workflow and up to five of your tools, and reviewed by a person before delivery.