

FREE FIELD GUIDE

# THE FORGE AGENT KIT

Build one useful AI agent every week - without wasting money on the wrong tools.

## WHAT THIS KIT GIVES YOU

A repeatable weekly build loop, a tool-budget guardrail, three copy-ready agent blueprints, and a seven-day shipping checklist.

Version 1.0 | [theforgearsenal.com/agent-kit](https://theforgearsenal.com/agent-kit)

No income promise. No giant software stack. Start with one recurring task and prove that the agent works.



# The weekly agent build loop

The fastest way to become good at agents is to ship one small, useful workflow at a time. Each build should have one user, one recurring task, one measurable output, and one clear failure condition.

| DAY | ACTION                    | DONE WHEN                                              |
|-----|---------------------------|--------------------------------------------------------|
| 1   | Choose one recurring task | You can describe the input and output in one sentence. |
| 2   | Write the manual workflow | A human could follow the steps without an AI tool.     |
| 3   | Build the smallest agent  | One model and only the tools the task truly needs.     |
| 4   | Run five real tests       | Failures are recorded instead of hidden.               |
| 5   | Add one guardrail         | Risky or low-confidence outputs require review.        |
| 6   | Package the result        | You have a prompt, checklist, and short demo.          |
| 7   | Publish the proof         | Show the input, output, limits, and next improvement.  |

## THE RULE

Do not add a second model, framework, database, or paid plan until a real test proves the first version cannot do the job.



# Agent opportunity scorecard

Score each candidate task from 0 to 2. Start with the highest-scoring task that has low data risk.

| QUESTION                        | 0                | 1               | 2               |
|---------------------------------|------------------|-----------------|-----------------|
| How often does it happen?       | Rarely           | Monthly         | Weekly or daily |
| How clear is the output?        | Subjective       | Partly defined  | Easy to verify  |
| How repetitive are the steps?   | Always different | Some repetition | Mostly repeated |
| How much time could it save?    | < 15 min         | 15-60 min       | > 60 min        |
| Can a person review the result? | No               | Sometimes       | Always          |
| How sensitive is the data?      | High risk        | Moderate        | Low risk        |

## GREEN LIGHT

A score of 9-12 is a strong first build. A score below 7 usually means the task is too rare, unclear, or risky for a beginner agent.

**Good starter jobs:** summarize recurring research, repurpose an approved piece of content, classify inbound leads, turn meeting notes into assigned tasks, or prepare a daily operations brief.

**Avoid at first:** autonomous purchases, financial decisions, medical or legal advice, irreversible account changes, or anything that handles credentials without a secure approval step.

# Blueprint A: Daily research brief

## OUTCOME

Turn a defined source list into a short morning brief with links, dates, changes, and suggested follow-up questions.

## Inputs

- A fixed list of trustworthy sources or URLs.
- The topic, audience, and maximum brief length.
- A delivery destination such as email, Telegram, Notion, or a file.

## Workflow

- Collect only new items since the previous run.
- Remove duplicates and separate facts from interpretation.
- Summarize each item with its source and publication date.
- Flag uncertainty instead of inventing missing information.
- Send the brief and retain a simple run log.

## Copy-ready starting prompt

You are a research brief agent for [AUDIENCE]. Review only the supplied sources for updates about [TOPIC] since [DATE/TIME]. Return: (1) the five most important changes, (2) one-sentence impact for each, (3) source link and date, (4) unanswered questions. If a claim cannot be verified from a supplied source, label it UNVERIFIED. Maximum 500 words.

**Starter budget:** \$0 first. Use a model and automation tier you already have; pay only after five useful runs.

# Blueprint B: Content repurposer

## OUTCOME

Turn one approved video transcript into a useful newsletter outline, three short posts, and five clip ideas while preserving the creator's claims and tone.

## Inputs

- A final transcript or article that you own or have permission to use.
- The target audience and one call to action.
- Examples of your real tone and phrases to avoid.

## Workflow

- Extract the central problem, proof, process, and takeaway.
- Quote only text that appears in the source.
- Create platform-specific drafts instead of copying one format everywhere.
- Run a claim check against the original transcript.
- Require human approval before anything is published.

## Copy-ready starting prompt

```
Repurpose the supplied transcript for [AUDIENCE]. Create: one newsletter outline, three short posts, and five 30-60 second clip concepts. Keep every factual claim grounded in the transcript. Do not add income claims, customer results, statistics, or quotations that are not present. End each draft with this CTA: [CTA]. Add a CHECK line listing any claim that needs human verification.
```

**Starter budget:** \$0-\$20/month until the workflow consistently saves more time than it costs.

# Blueprint C: Lead qualifier

## OUTCOME

Turn an inbound form submission into a structured summary, fit score, missing questions, and a draft response for human review.

## Inputs

- A clear description of the offer and who it is for.
- Qualification rules: problem, urgency, budget range, authority, and fit.
- An approved response template and escalation rules.

## Workflow

- Validate required fields and reject obvious spam.
- Summarize the stated problem without guessing.
- Score against written criteria and show the reason for each score.
- Draft follow-up questions for missing information.
- Send to a human for approval; never auto-close or auto-reject at first.

## Copy-ready starting prompt

Review this inbound lead for [OFFER]. Use only the submitted information. Return: problem summary, desired outcome, urgency, budget signal, decision-maker signal, fit score from 0-10 with reasons, missing questions, and a concise draft reply. Never infer protected traits or personal information. If the evidence is incomplete, say INSUFFICIENT INFORMATION and ask a question.

**Starter budget:** \$0 first. A form, spreadsheet, and existing model are enough to prove the workflow.



# Wrong-tool guardrails

A bigger stack is not a better system. These rules protect your cash and keep the build understandable.

- **Free-first rule:** prove the workflow with free tiers or tools you already own.
- **Monthly-only rule:** avoid annual plans until the agent has delivered value for four consecutive weeks.
- **One-model rule:** use one primary model until a documented failure requires another.
- **Human checkpoint:** keep approval before sending, publishing, purchasing, deleting, or changing accounts.
- **Data rule:** do not paste passwords, recovery codes, payment data, or private customer records into a casual prototype.
- **Exit rule:** every paid tool needs a cancellation date, export path, and replacement plan.

## The 10-minute tool test

- Can it accept your real input format?
- Can you inspect and correct the output?
- Can it connect to your existing workflow without a fragile workaround?
- Can you export your data and prompts?
- Does the monthly cost make sense before revenue or verified time savings?

### BUY ONLY AFTER PROOF

A tool passes when it solves the same real task five times with acceptable quality. A flashy demo is not proof.

# Your seven-day build checklist

- I chose one recurring task with a clear input and output.
- I wrote the manual steps before adding tools.
- I set a weekly cost ceiling.
- I built the smallest version with one primary model.
- I tested five real examples and saved the failures.
- I added a human approval step for risky actions.
- I recorded a short demo showing input, output, limits, and next improvement.
- I shared the build and selected next week's single improvement.

## BUILD WITH MIKE

Want help turning one workflow into a working agent? Visit the Forge Arsenal for the live kit, weekly demos, and the Founding Agent Sprint.

[theforgearsenal.com/agent-kit](https://theforgearsenal.com/agent-kit)

[youtube.com/@Mikesaiforge](https://youtube.com/@Mikesaiforge)

This guide is educational. Results depend on the task, data, tools, testing, and implementation.  
Review important outputs before acting.