AI AGENTS · PDF TEMPLATE

AI Agent Build Brief

A working business agent is mostly instructions, context and tools; the model is the easy part. Fill this in before anyone builds.

6 SECTIONS · 13 CHECKS · FILL-IN TEMPLATE



01 One job and what done looks like

FILL IN

The job in one sentence

Owner of the agent

Example of excellent output 1

Example 2

Example 3

02 Instructions a new hire could follow

FILL IN

The role: who it works for, what it is responsible for

The steps: the usual order, and where it may deviate

The standard: what a good result contains

The limits: what it must never do, when to stop and ask

Be specific. "Be helpful and accurate" adds nothing. "Quote prices only from the current price list, and say so if a product is not on it" changes behavior.

03 The right context, and only that

FILL IN

SOURCE THE TASK USES	LIVE LINK OR COPY?	OWNER

› **Relevant beats complete.** Give the agent what the task uses, not every document you own.

› **Current beats convenient.** Link to the live source so it does not quote last year's pricing.

04 Only the tools the job needs

FILL IN

TOOL NAME	WHEN TO USE IT	READ OR WRITE	NEEDS APPROVAL?

- **Fewer, task-shaped tools:** one that checks availability and books beats two raw ones.
- **Readable results:** names and plain fields, not internal IDs; filter large results.
- **Clear, distinct names:** prefix by system, such as crm_search and helpdesk_search.
- **Useful errors:** "No customer found; try searching by email" helps it recover.

05 Guardrails before anyone relies on it

5 CHECKS

- **Narrow permissions** for every tool, using each user's own access where possible.
- **A clear way to escalate:** when unsure, the agent stops and asks.
- **A stopping condition,** such as a maximum number of steps per run.
- **Approval steps** for anything that sends, pays, commits or deletes.
- **Logs of every run,** reviewed during the first weeks.

06 Test, launch small, then keep it current

SIGN OFF

- **Test** on your three examples plus the awkward cases: missing information, conflicting sources, out-of-scope requests.
- **Launch** to one person or one queue, compare with the old process, widen only when the numbers hold.

AFTER LAUNCH	WHO / WHEN
Owner of the instructions	
Current version number	
Log review schedule	
Rerun the tests when instructions, tools, context or model change	

NEXT STEP

AI agent services

www.doitwithfriday.com/services/agents-skills/

Build with Friday →

BASED ON FRIDAY FIELD NOTES

www.doitwithfriday.com/blog/how-to-build-an-ai-agent/

www.doitwithfriday.com/blog/how-to-test-ai-agents/

General guidance, not legal or security advice for your specific setup. Questions: hello@doitwithfriday.com