

S-10

Weekly Sprint Report template

The end-of-week report template with criteria status, pull requests, measurements, hours and risks, plus a filled example.

Axdale. Software engineering by Michael Lip, Warsaw, Poland (CET).

Website <https://www.axdale.com>. Code portfolio on GitHub at <https://github.com/theluckystrike>.

Version 1.0. 23 October 2026.

You get this report by 18:00 CET on the last Working Day of each Week. Part A is the blank template. Part B is a filled example for a fictional client, and it follows the Statement of Work example in the same document set.

Part A. Template

1. Reference

Field	Details
Client	_____
SOW number	_____
Week	Free Week (week 1) / Paid Week number _____
Working Days	_____ to _____
Hours used	_____ of 40
Report date	_____

2. Summary

Three sentences at most. What shipped, what it means for your users and what is next.

3. Acceptance criteria

No.	Criterion in short	Status	Evidence
AC1		Met / Not met / Partly met	
AC2			
AC3			

4. Pull requests

PR	What it does	State	Criterion
		Merged / In review / Draft	

5. Measurements

Measure	Last Week	This Week	Target

6. Hours

Day	Hours	Main work
Total		of 40

7. Blocks and time lost

8. Production actions

Action	Approved by and when	Result
None /		

9. AI providers used

10. Risks

Risk	Likelihood	Effect	What I do about it
	Low / Medium / High		

11. Decisions needed from you

1. _____

12. Plan for next Week

1. _____

2. _____

Next Week estimate _____ hours at the agreed hourly rate. Invoice sent on _____.

Part B. Filled example (fictional)

The client, people and numbers below are fictional. They show how a real report reads.

1. Reference

Field	Details
Client	Example Client Ltd
SOW number	EXAMPLE-02
Week	Paid Week 2
Working Days	9, 10, 12, 13 and 16 November 2026 (11 November is a public holiday in Poland)
Hours used	38.5 of 40
Report date	Monday 16 November 2026

2. Summary

The reply drafter now runs against 40 real past tickets and passes the rubric on 36. The evaluation is a required CI check, so a change that lowers quality can't merge. Next Week adds staff feedback buttons and a staged production rollout behind the flag, with your written approval for each deploy.

3. Acceptance criteria

No.	Criterion in short	Status	Evidence
AC1	Evaluation set grown to 40 anonymized cases	Met	PR #41, file evals/drafts/cases.jsonl
AC2	At least 34 of 40 cases pass the rubric	Met, 36 of 40	CI run on PR #44, report attached
AC3	Evaluation is a required CI check on the main branch	Met	Branch protection setting, PR #43

No.	Criterion in short	Status	Evidence
AC4	Each request stops at 4,000 output tokens and the daily spend stops at USD 20 on staging	Met	Tests <code>drafts.caps.spec.ts</code> , staging log
AC5	p95 latency under 15 seconds on the evaluation set	Partly met, 16.2 seconds	See risk 1 and decision 1

4. Pull requests

PR	What it does	State	Criterion
#41	Grows the evaluation set from 15 to 40 cases	Merged	AC1
#42	Fixes timeouts on tickets longer than 3,000 words	Merged	AC5
#43	Adds the evaluation as a required CI check	Merged	AC3
#44	Rubric update and report output in the CI summary	Merged	AC2
#45	Token cap per request and daily spend cap	In review	AC4

5. Measurements

Measure	Week 1	This Week	Target
Evaluation cases	15	40	40
Cases that pass the rubric	12 of 15 (80%)	36 of 40 (90%)	85% or more
Drafts with at least 1 help article link	14 of 15	39 of 40	All
p95 latency	13.8 seconds	16.2 seconds	Under 15 seconds
Model cost per draft on staging	USD 0.021	USD 0.017	Under USD 0.03
Prompt injection cases blocked	3 of 3	6 of 6	All

6. Hours

Day	Hours	Main work
Mon 9 Nov	8.0	Selected and anonymized 25 new cases with the support lead's export
Tue 10 Nov	7.5	PR #41 and PR #42
Thu 12 Nov	8.0	PR #43, CI check and branch protection
Fri 13 Nov	7.5	PR #44, rubric review
Mon 16 Nov	7.5	PR #45, latency tests, this report

Day	Hours	Main work
Total	38.5 of 40	

7. Blocks and time lost

The staging test key hit its spend cap on Tuesday at 14:10 CET. Jane Doe raised the cap at 15:30 CET, and I worked on PR #42 in the meantime, so I lost no time.

8. Production actions

Action	Approved by and when	Result
None	Not applicable	Not applicable

9. AI providers used

Anthropic only, as approved in the Schedule. No other provider saw the Client's code or tickets.

10. Risks

Risk	Likelihood	Effect	What I do about it
p95 latency is 16.2 seconds, above the 15 second target	High	Staff wait longer on long tickets	Stream the draft so the first words show in about 2 seconds, or allow 20 seconds for tickets over 3,000 words (decision 1)
Help articles older than 12 months give outdated answers	Medium	Some drafts cite old steps	Add the article date to the draft and flag articles older than 12 months

11. Decisions needed from you

1. Choose one. Stream the draft (my default, about 6 hours of work), or accept 20 seconds for tickets over 3,000 words.
2. Name 2 support staff who will try the drafts on staging next Week and leave written feedback in the tool.

12. Plan for next Week

1. Staff feedback buttons on each draft, stored on staging.
2. Streaming of drafts, if you choose it in decision 1.
3. A cost and quality page for the support lead.
4. A staged production rollout plan behind the flag. Each deploy needs your written approval.

Next Week estimate 40 hours at the agreed hourly rate. Invoice sent on 16 November 2026.

Axdale. Michael Lip, independent engineer, Warsaw, Poland (CET). michael@axdale.com.

Axdale, Michael Lip. Warsaw, Poland (CET). michael@axdale.com.
www.axdale.com. Code record github.com/theluckystrike. Version
1.0. 23 October 2026.