

The Agent Standup

Ed Schaefer · Version 1.0 · September 2026 · edward-schaefer.com/templates/agent-standup · Free to share

Loops keep running whether or not you're watching. This page is how you check in on them: the same practice as a team standup, run for a fleet of agents. What happened, what's running, what's stuck, what needs you, what's next. The work doesn't stop when this page is generated, and it doesn't wait for you to read it. An agent here means any AI worker running a task on its own: a coding agent, a drafting worker, or a loop that dispatches other agents.

Cadence: every day. Shrink it on quiet days; most sections below compress to one line. Never skip a day. A skipped day breaks the one thing this is for: seeing today against yesterday against the day before.

How to read this page

- **Stalled:** the item is still assigned or running, but nothing has moved past the age threshold you set below. It hasn't failed; it has gone quiet while still alive.
- **Silent:** no signal at all, no heartbeat, no error. Silence is its own status here, never a blank line.
- **[CLAIMED]:** a self-report. The agent, or the system relaying it, said so. Worth reading; never sufficient on its own.
- **[VERIFIED]:** checkable by something other than the thing making the claim: a passing test, an exit code, a diff, a second reader.
- **Evidence link:** the actual artifact behind a "done" claim, such as a branch or commit SHA (the pointer into version control anyone can open), a test run, or a captured output. A restated summary does not count.
- **Heartbeat:** the footer that proves this page itself is current: when it was generated, what window it covers, how many agents it could reach.
- **[MACHINE]:** an existing or buildable feed could fill this line without a person typing it. **[HUMAN]:** nothing computes this; someone states it. Today, fill in every line by hand regardless of its tag; the tag marks what a later version could automate, and nothing is skippable now.

0. One-line summary

[MACHINE] [VERIFIED]

Write this line last, after every section below; it is the roll-up.

Summary: _____

1. Needs you today

[HUMAN] budget · [MACHINE] [VERIFIED] count · [MACHINE] [CLAIMED] ranking

Your decision budget for this morning: _____ (default: **3**). Rank items oldest-first, same as a board-walk; list up to the budget, then stop. Anything stalled past the threshold in section 2, or blocked for more than a day in section 3, escalates here and counts against the budget.

#	Item	Why it needs you	Claimed / Verified
1			
2			
3			

Falls to tomorrow (past the budget): _____ (or "none")

Track your actual count for a week before you change the default of 3; let the number come from what actually shows up.

2. Stalled or silent

[MACHINE] [VERIFIED]: elapsed time since the last real update, measured from timestamps.

Render "None" if nothing qualifies. Default staleness threshold if you have no data of your own yet: **2–3 days**; once you know your own cycle time, use that instead.

Item	Last activity	Stalled / Silent	Owner or next step

3. Blocked

[MACHINE] [VERIFIED] that something is blocked and which trigger fired · [CLAIMED] for any suggested next step.

Named every time; a blocked item that sits unremarked is the failure.

Item	Blocked on (named trigger)	Suggested next step [CLAIMED]	Owner

4. Failed / problems found

[MACHINE] [VERIFIED] pass/fail verdict · [CLAIMED] the agent's stated reason and its confidence line.

For each failure, the reason the agent gives is a claim; check it against the actual output before you act on it.

Agent	Item	Verdict	Agent's stated reason [CLAIMED]	Confidence (0.0 to 1.0) [CLAIMED]	Evidence
		FAIL			

Render "None" if nothing failed.

5. In progress

[MACHINE] [VERIFIED] that something is running · [CLAIMED] what it's actually doing.

Agent	On what	Running since	What it's doing [CLAIMED]

6. Done since yesterday

[MACHINE] that a branch, commit, or test run exists · [VERIFIED] wherever that evidence backs the claim · otherwise render it as **CLAIMED, self-reported**, never upgraded silently.

Agent	Item	Evidence (branch / SHA / test / output)	Verified / Claimed	Confidence [CLAIMED, never verified]

7. Not yet reviewed by you

[HUMAN]: nothing in a typical fleet's tooling computes a check's own scope yet, so this is hand-filled.

For EVERY item marked done above, including the ones marked claimed-only: what the check covered, and what it didn't. A claimed-only item with no check is the first line here, never the one skipped.

Item	Check covered	Check did not cover

8. Planned today / what to expect by end of day

[MACHINE] for what's queued · [HUMAN], explicitly a guess, for what to expect.

Queued: _____

What to expect by end of day (a guess): _____

9. Yesterday's directives, closed or not

[HUMAN]: most fleets have no log of what a person told the fleet to do, so read this back from your own notes until one exists.

You told the fleet	Closed / Not closed / Partly	Note

10. Heartbeat

[MACHINE] [VERIFIED]

Generated: _____ · Covers: _____ through _____ · Next check:
_____ · Agents confirmed alive: _____ of _____ · Spend since last page: \$_____ [MACHINE]
[VERIFIED] · Read time: ~_____ min

Worked examples

Both examples use generic placeholders only.

Example A, a quiet day

0. Summary: Fleet quiet: nothing stalled, nothing blocked, 2 shipped since yesterday.

1. Needs you today: Budget: 3. None. Falls to tomorrow: none.

2. Stalled or silent: None.

3. Blocked: None.

4. Failed / problems found: None.

5. In progress: Agent Ridge, on Project Aurora's export job, running since 06:10 [VERIFIED]; writing the nightly CSV [CLAIMED].

6. Done since yesterday: Agent Ridge, "Retry logic for the export job"; evidence: branch fix/export-retry, PASS on 14 tests [VERIFIED]; confidence 0.8 [CLAIMED]. Agent Ridge, "Update the README"; evidence: commit a1b2c3d [VERIFIED]; confidence 0.95 [CLAIMED].

7. Not yet reviewed: Retry logic: unit tests pass; no run against a real failed export yet. README update: commit exists; nobody has read the diff.

8. Planned today: Queued: one more pass on the export job's error path. Expect by EOD [HUMAN, a guess]: error-path pass done, nothing else queued.

9. Yesterday's directives: "Hold the export job at read-only until the retry fix lands": closed; retry fix shipped, read-only hold lifted.

10. Heartbeat: Generated 07:00 · Covers yesterday 07:00 through today 07:00 · Next check tomorrow 07:00 · Agents confirmed alive: 1 of 1 · Spend since last page: \$4 · Read time: ~1 min.

Example B, a bad day

0. Summary: Fleet: 1 stalled, 1 blocked, 1 failed, 2 shipped since yesterday (1 verified, 1 claimed). 3 decisions need you this morning.

1. Needs you today: Budget: 3. 1. Approve the rollback for Project Vesper's pricing job: [VERIFIED] the job wrote bad output while it kept running unattended; needs your call before it runs again. 2. Decide whether to keep retrying the flaky login test on Project Aurora: [CLAIMED] the agent's own read is that it is a timing issue, not corroborated yet. 3. Project Vesper's audit sync, escalated from section 2: stalled 3 days, no owner. Falls to tomorrow: none.

2. Stalled or silent: Project Vesper's audit sync: last activity 3 days ago, no error; STALLED, silent. Owner: unassigned. Escalated to section 1.

3. Blocked: Project Aurora's deploy step: blocked on missing API credential (named trigger: credential_missing). Suggested next step [CLAIMED]: rotate the key. Owner: you.

4. Failed / problems found: Agent Vale, Project Vesper's pricing job: FAIL. Agent's stated reason [CLAIMED]: "input schema changed upstream." Confidence: 0.4 [CLAIMED]. Evidence: run log exit code 1, output tail attached.

5. In progress: Agent Vale, on Project Vesper's rollback, running since 06:45 [VERIFIED]; reverting the pricing job to the last good commit [CLAIMED].

6. Done since yesterday: Agent Vale, "Add logging to the pricing job"; evidence: branch feat/pricing-logs, PASS on 9 tests [VERIFIED]; confidence 0.85 [CLAIMED]. Agent Moss, "Draft the credential rotation runbook"; no evidence link attached; CLAIMED, self-reported.

7. Not yet reviewed: Credential rotation runbook: no evidence link, nothing checked; treat as not done until a file exists. Pricing-job logging: unit tests pass; never run against yesterday's actual bad input.

8. Planned today: Queued: rollback completion, credential rotation. Expect by EOD [HUMAN, a guess]: pricing job back on last good commit; deploy still blocked unless the credential is rotated today.

9. Yesterday's directives: "Stop the pricing job the moment output looks wrong": not closed; it ran twice more before anyone stopped it, and no kill control existed for that job. Flagged as a gap.

10. Heartbeat: Generated 07:00 · Covers yesterday 07:00 through today 07:00 · Next check tomorrow 07:00 · Agents confirmed alive: 2 of 3 · Spend since last page: \$31 · Read time: ~3 min.