PAPER 3 NOVEMBER 2026 BY THE LAB 5 MIN READ VERSION 1.1

How laundry keeps your Git tidy

Laundry ran 52 times in 15 of the lab's repositories over two months, archived and removed 196 branches, left 36 more archived in place, and refused 403. A check on 2026-11-02 found an archive for all 196. The refusals are the finding.

52 runs. 196 branches. All 196 archives found.

52	196	36	403
runs, one run of laundry each	branches archived, then removed	archived, left in place	branch refusals

Why we built it

Every developer knows the pause before `git branch -D`. Is this the branch that merged, or the one with Tuesday's fix? Most of us answer by leaving it. Then there are hundreds.

A lab that works with agents makes branches faster than a person can read them. Every clone, and every agent that reads `git branch -a`, pays for them.

Where other tools stop

`git branch --merged` lists a branch only when its tip sits inside the branch you are on. A squash merge writes a new commit, so a branch that has landed still looks unmerged, and a careful person keeps it. GitHub can delete a pull request's branch when it merges, but only on GitHub; every clone keeps its copy. A script that pipes a list into `git branch -D` deletes on a guess, and a wrong guess loses work nobody can get back.

Why safety came first

A lost commit is an afternoon of someone's thinking. We wanted a tool that is safe to hand to an agent, which meant it had to be safe to be wrong. It proves before it removes and archives before it deletes. When it cannot prove, it refuses and says why. On its own it only surveys and reports. A person says yes before `--apply` removes anything. This paper is its record from 26 August to 25 October 2026, counted by the tool itself.

What it saves

A person facing hundreds of branches reads each name and decides, one at a time, what is safe. Laundry does that reasoning once, writes it down, and hands back only the branches it refused, each with its reason.

An agent pays the same cost in tokens. We measured it once, in an agent trial named AB1 on 2026-10-30. An agent tidied a seeded repository, once by hand with git and once with laundry. There were 36 runs in 6 cells, 3 runs per cell, so the sample is small. They used 36,399,339 tokens in all. With laundry the agent was right in 18 of 18 runs. By hand it was right in 15 of 18.

In 4 of the 6 cells laundry used 12.0% to 44.6% fewer tokens. In the other 2 it used 6.9% and 0.8% more.

What laundry does

Laundry surveys every local and remote branch and sorts each one: merged by ancestry, merged by a pull request, squash-merged, superseded, held by a worktree, protected, current, open on a pull request, or unmerged. For each branch it can prove safe, it writes a snapshot ref, verifies the ref points at the exact commit, and only then removes the branch. Anything it cannot prove, it refuses.

Every removal is archived first. One command restores it to the exact commit: `--restore <run>` brings back everything a run archived, `--restore-branch <name>` one branch. There is no model. It reads git, refreshes your remote's branch list with `git fetch --prune`, and asks `gh` when it is present. Nothing of ours is called.

Every run writes one row to the repository's own git directory: when, which version, how many branches it examined, planned and refused, and why. That row is the data in this paper.

What the record shows

52 runs in 15 repositories, 2026-08-26 to 2026-10-25, across 8 versions from 1.0.1 to 1.7.0-dev.5. The tool examined 618 branches. It planned 232 removals, wrote a snapshot ref for every one, and completed 196. The other 36 were archived and left in place. It refused 403 branches.

Use grew with trust. August had 8 runs and 33 removals. September had 10 and 27. October had 34 runs and 136 removals, more than the two months before it together.

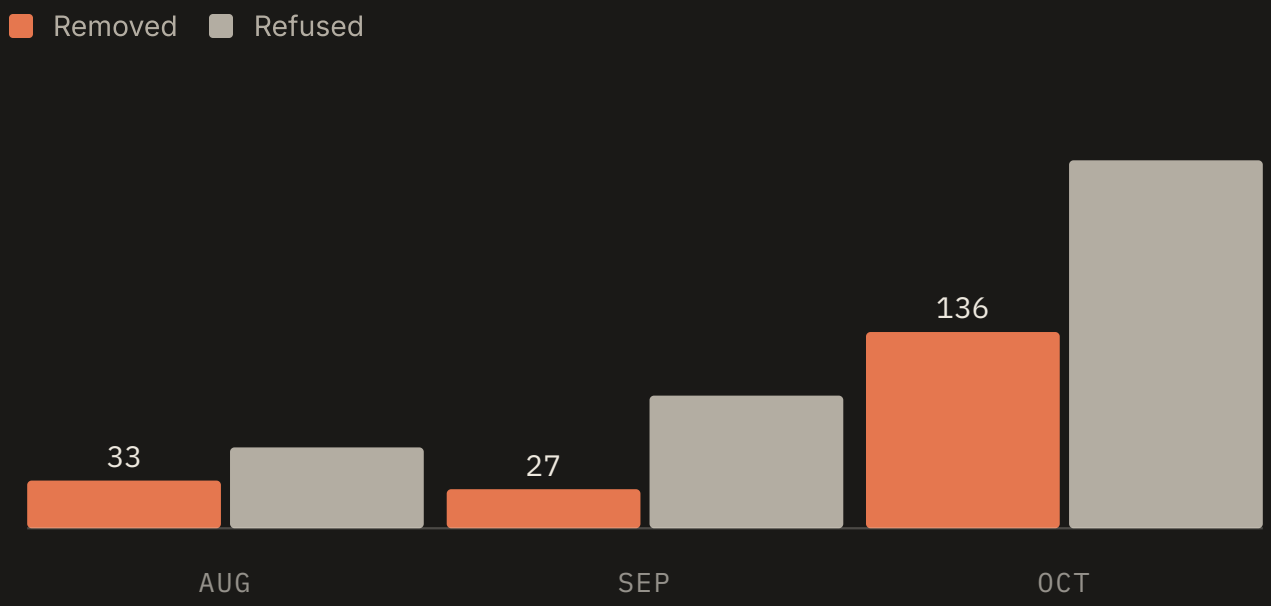


Figure 1. Removed and refused, by month.

MONTH	RUNS	EXAMINED	REMOVED	REFUSED	SECONDS
August	8	77	33	56	107
September	10	119	27	92	100
October	34	422	136	255	353

The refusals are the finding. Of 403, the reasons were: unmerged work 102, a protected branch 99, the head of an open pull request 66, a pull request marked merged that no proof confirmed 59, a branch held by a worktree 45, the current branch 12, and the rest under a goal's scope or a keep rule. A tired person, or an eager agent, would have deleted each of those. Laundry left every one in place and printed the reason.

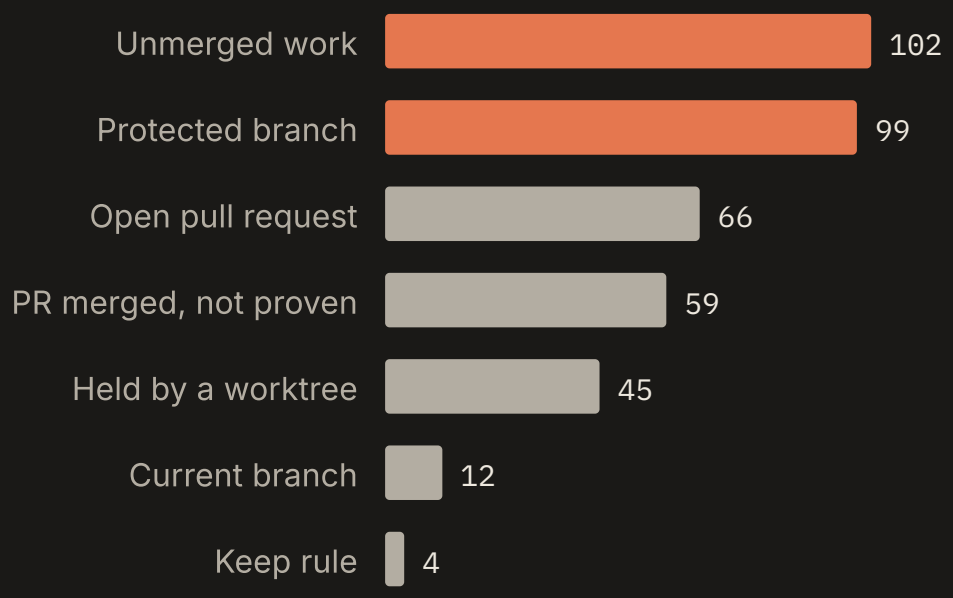


Figure 2. 403 refusals, by the reason Laundry printed. Each one is a branch left in place.

45 of 52 runs completed. 3 ended partial, 3 failed, 1 stopped for a human. All 33 failed steps share one cause: git refused the delete because it could not see the branch as merged, and laundry did not force it. The snapshot ref had already been written each time. One run on 25 October did this 22 times and ended partial, saying so.

The meter's own counter for lost commits is a fixed zero, so we do not quote it. We checked instead. On 2026-11-02, 196 of 196 removed branches still had an archive ref at the exact commit the log recorded. That check reads today's refs, so it checks the record. It does not prove the future.

Since this record

The meter holds nothing after 25 October. Three things happened since, in the lab's own record.

An adversary pass, ADV1, ran on 2026-10-30. One agent wrote 21 attack scripts. None made `--apply` lose a commit. One found that a survey could drop an archive

ref, since fixed. Version 1.7.3 closed four smaller findings.

A team trial, TEAM5, had 5 people run `--apply` on one repository at once. Across 24 trials over 8 stages, no work was lost.

On 2026-11-02 we re-ran the suites on the 1.7.0 tag. The smoke suite passed 199 cases, failed 0 and skipped 3 with printed reasons. 25 invariant tests passed.

What we learned

A refusal is worth more than a removal. The 403 refusals make laundry safe to hand to an agent: the agent surveys, and a person says yes to `--apply`.

The meter needs reading. The first draft of this paper said 232 branches were removed, because the meter's counter counts removals planned. The step log showed 36 ended archived and in place. The first version also said no commit was lost, resting on a constant. The 196 of 196 check replaced it.

Failing a step beats forcing it. Every failed step in two months was git saying no. Versions 1.7.0 to 1.7.3 shipped exact proofs for cases where the refusal was git's caution and not a real divergence.

Limits

One user, one lab. Every run was in the lab's own repositories, by the founder or by agents working for him. No outside team has run laundry yet.

The meter holds no rows after 2026-10-25 and none for versions 1.7.1 to 1.7.3. Those versions rest on the tests and trials named above, not on the meter. This paper was built from the 1.7.3 line's record, before its tag.

The meter cannot see the agent that invoked the tool; the agent trial measured that once, with 3 runs per cell. The archive check reads today's refs and covers commits

only. Examined, removed and refused are separate counters and are not expected to sum.

Try it

Laundry is at household.dev/tools/laundry. Clone the repository, run it with no flags to see the survey, and read the refusals before anything else. It is Apache-2.0.

Questions about one line of this record go to hello@household.dev.

Try Laundry

Nothing lost in the wash.

Cite this

```
Household Intelligence Labs Inc. How laundry keeps your Git tidy. www.household.dev/research/laundry, 3 November 2026.
```

Questions

hello@household.dev

Earlier versions

Version 1.0, 27 October 2026

Sources

Counted by the tool itself, one row per run, read from every repository on the founder's machine. Exact on this page. Nothing here is reported by an agent about itself.

The data

Exact counts in data/tool-usage/laundry.json: One JSON row per run, written by the tool itself to its repository's own git directory, read from every repository on the founder's machine.

On the 1.7.0 tag, 2026-11-02, the smoke suite passed 199 cases, failed 0 and skipped 3 with printed reasons, and 25 invariant tests passed

laundry lab, reports/1.7.0-claims-settled.md, C13 (evidence/1.7.0/tag-smoke.txt and tag-invariants.txt)

On 2026-11-02, 196 of 196 removed branches still had an archive ref at the exact commit the log recorded; the check reads today's refs in the lab's own repositories and covers commits only

laundry lab, reports/1.7.0.md, section Commits lost (freeze_meter.py), and reports/1.7.0-claims-settled.md, C1

22 branches in one run were archived and left in place

one repository's .git/laundry/log.jsonl, run 2026-10-25T0916Z-6fdd

33 failed steps across 6 runs, one cause

every .git/laundry/log.jsonl on the founder's machine, read 2026-10-27, recomputed 2026-11-02 in laundry lab reports/1.7.0-claims-settled.md

8 versions of the tool

data/tool-usage/laundry.json, versions

git branch --merged lists a branch only when its tip is reachable from the branch you are on, so a squash-merged branch looks unmerged

git-scm.com/docs/git-branch (--merged), read 2026-10-28

GitHub's automatic branch deletion removes the head branch on GitHub after a merge; local clones keep theirs

docs.github.com, Managing the automatic deletion of branches, read 2026-10-28

Agent trial AB1 (2026-10-30): 36 runs used 36,399,339 tokens in 6 cells, 3 runs per cell; the agent was right in 18 of 18 runs with laundry and 15 of 18 by hand; no work was lost in either arm

laundry lab, reports/1.7.0.md, section The agent trial, and experiments/EXPERIMENTS.md, row AB1

AB1, tokens per run over the correct runs: laundry used 12.0%, 28.4%, 34.5% and 44.6% fewer in four cells, and 6.9% and 0.8% more in the other two; at 182 branches it was slower in wall time

`laundry lab, reports/1.7.0.md, the AB1 table, and reports/1.7.0-claims-settled.md, C17 and the mock's claims`

Adversary pass ADV1 (2026-10-30): one agent, 21 attack scripts; no lost commit from --apply; a survey could drop an archive ref, since fixed

`laundry lab, experiments/EXPERIMENTS.md, rows ADV1 and ADV1 (correction), and reports/1.7.0-claims-settled.md, C2`

Version 1.7.3 closed the four smaller ADV1 findings

`laundry 1.7.3 build record (docs/releases/v1.7.3/summary.md), section Genesis`

Team trial TEAM5 (2026-10-30): 5 people ran --apply on one repository at once; 24 trials over 8 stages; 0 work lost

`laundry lab, experiments/EXPERIMENTS.md, row TEAM5, and reports/1.7.0.md, section The team trial`

Questions about one line of this paper go to hello@household.dev.