



- 1 **Context: purge a finished run (admin only)** The harder of the two ways to clear a run, shown by the flame on a run's row (**Experiments → runs pane**) and by **Run detail → Purge run** — the same dialog from either place. **Cleanup** reclaims the bulk and *keeps* the `experiment_run` row and its `run_result` scorecard; **Purge** drops the run row, which cascades the scorecard away, so the run leaves the comparison. Both are archive-first by default (the archive can be declined, with a warning), and both are refused on a running or paused run. Purging removes the run itself, along with its scorecard — the comparison loses this run entirely.

---

- 2 **Archive first — the default, not a wall** After a purge the archive may be all that survives of the run, so it is ticked by default and written to durable storage *before* the delete, its object key recorded in the `audit_log` entry. **If the write fails, nothing is deleted** — the dialog says so and offers Retry. Unticking is allowed, because purging a bad run should not force an archive nobody wants, but it is a warned choice: the acknowledgment below changes to “this run cannot be recovered in any form,” and the audit entry records that no archive was taken (ALGORITHMS §13).

---

- 3 **What a purge removes** Everything a cleanup removes — `run_enrollment`, `chat_round` and its candidates, message and its recipients, the run's `audit_log` rows — **plus** the `experiment_run` row itself, whose delete cascades `run_result` away so no scorecard is ever left dangling.

---

- 4 **Why purge exists** A definition is held **immutable while any run references it** (a database RESTRICT rule), so a study, nudge, model configuration or cohort cannot be edited or deleted until its last run is purged. Cleanup does not release that lock, because it keeps the run row.

---

- 5 Switch to the gentler option, which keeps this run's scorecard.

---

- 6 Cancel without clearing anything.

---

- 7 Permanently remove this run and its scorecard.