

AI*Right: what's still open (plain English), team version

For the AI*Right team · Oz · Mon Oct 19, 2026, 4:05 PM ET (EOD update) · Source: gap-register.md + open-items-for-team.md

Who's who: Ed = Scale's AI engineering agent (a bot that builds and tests). Pam = Scale's AI project-management agent (a bot that plans and witnesses test runs). Oz = Scale's AI coordinator agent. Sal = Scale's AI sales-scoreboard agent. Jon = the only human approver of writes.

Closed or decided today: 14. 9 closed: G9, G10, G16, G19, G25, G30, G32, G44, and G39's gateway half. 4 decided: G12 gateway, G31 console, G33 Clef, G45 approvers (design done, build open). G29, the sign-off rule, has been applied. **Listed here: 33 open register gaps (including new connector gaps G46-G52), plus 8 open tasks from the team list.**

New since 3:07 PM:

- **Licensing resolved:** Scale has Entra ID P2 (15 seats, 14 assigned), so Conditional Access with MFA is licensed. Whether Jon's own account holds a P2 seat is UNKNOWN; Ed checks during setup.
- **Approver prototype:** approval-screen triage built (summary line, red-flag snippets, injection banner, isolation banner never green by default, isolation proof via context-set hash + tenant_id, kill button, 8 s approve delay, 8-per-session cap + 15-min lockout, trap item). 35/35 tests, synthetic, unwitnessed; all timings EXAMPLE.
- **G45 approvers DECIDED (option A):** per-client approvers in each client's own Entra tenant; Jon is sole approver for the pilot only, through Nov 30. Isolation that can't be proven is a hard block, with no override. No second-approver thresholds in the pilot. Design done; build open.
- **Pilot exit criteria EX-1 to EX-4 drafted:** audit trail from intake via connector, tenant isolation, wrong-write undo, judge catch on the 49-attack set. Jon is the EX-1 sign-off. Thresholds set after Thu Oct 22.
- **R0 v1.1 + live Clef calibration:** 62 synthetic docs (13 clean, 49 attacks). 183/183 HTTP 200, 0 writes, AUC 0.987, 0 false allows at 0.80/0.30 (and at the proposed 0.70/0.30). Unwitnessed; Pam witnesses Thursday's run before any threshold change.

1. Needs team input by Fri Oct 23 (14)

Item	What it is / why it matters	Owner	Due
G34 Gateway ID	How the results store proves a request really came from our gateway. The test version uses a shared password.	Ed	Fri Oct 23
G35 Key care	Who holds those keys and how often they're changed.	Ed	Fri Oct 23
G36 Network fence	Network rules so only the gateway can reach the results store, per client.	Ed	Fri Oct 23
G37 Approval service	The real system that records Jon's approvals; today it's a test file.	Ed	Fri Oct 23
G38 Refusal log	Blocked attempts are logged, but not yet in the standard audit format.	Ed	Fri Oct 23
G40 Client sign-in	How we tell which client a user belongs to (their own Microsoft, a guest login, or Okta). Getting it wrong could mix clients' data.	Ed; Jon	Fri Oct 23 (decide in Assess)
G41 Usage limits	Limits on simultaneous requests and bursts per client; the numbers are UNKNOWN.	Ed	Fri Oct 23
G42 Client isolation	Separate networks, keys and storage for each client, and where the data lives.	Ed	Fri Oct 23
Approval screen	The simple screen Jon uses to approve pilot writes. LiteLLM's built-in UI is ruled out. Ed's	Ed + Pam; team	Review Tue/Wed Oct

Item	What it is / why it matters	Owner	Due
	prototype is built (35/35 tests, synthetic): red flags, isolation proof, kill button, approve delay, session cap, trap item. Build vs buy for the real approval service (G37) is Jon's call Tue Oct 20. Please review the prototype Tue/Wed.	reviews	20-21
Pricing removal	How supplier prices are kept out of the pilot: by picking documents or by blacking out prices.	Pam + Ed	Fri Oct 23
Pilot go-live date	The go-live day within Nov 3-30, plus Pam's final test date (at most 5 business days before).	Pam	UNKNOWN
Accuracy bar	How many results get checked by hand, and what score counts as good.	Ed + Pam	UNKNOWN (after first baseline)
Clef tuning	First live calibration done (183/183 calls, AUC 0.987, 0 false allows at 0.80/0.30). Proposed change to 0.70/0.30 lets more clean docs through; it needs Pam's witness Thursday, then Jon's signature.	Ed; Pam witnesses	Thu Oct 22 test set
Clef unknowns	Speed, rate limits, error format, and how long Cloudflare keeps what we send are not documented.	Ed	Fri Oct 23

2. Ed or Pam working (10 lines, covering 18 gaps)

Item	What it is / why it matters	Owner	Due
G11 Booby-trapped documents	Our top pilot risk: hidden text inside a document trying to trick the AI. A 40-document attack test runs Thursday.	Ed; Pam witnesses	Thu Oct 22, 2:00 PM ET
G1, G4 Pilot building blocks	Reusable templates with the approval step built in, plus the first written rules (which tools are allowed, approval before any write).	Ed	Before pilot; exact date UNKNOWN
G3 Audit records	Log format is set; still to define the "evidence pack" for auditors. Which audit standard applies is UNKNOWN.	Pam; Ed	UNKNOWN
Clef live test	Ran today (2:51 PM ET) and came back green (HTTP 200): Clef said "proceed" on a test document, and the write was still held for human approval, with nothing written. The first try at 2:38 PM failed safely because of a bad account setting. Cost so far is under a tenth of a cent of the \$10 cap. Pam still needs to witness it.	Ed; Pam witnesses	Done today; witness date UNKNOWN

Item	What it is / why it matters	Owner	Due
Rehearsal run	Practice run in the real environment; waits on the go-live environment decision (due Mon Nov 2).	Ed; Pam witnesses	UNKNOWN
G14, G15 Plan hygiene	Stop each sale becoming a custom build, and line up the two roadmaps (6-phase vs 90-day).	Pam	Fri Oct 23
G17, G20 Claims	Say only what's true about "sovereign" AI, and keep a list of ChatGPT-style features we won't build.	Pam	UNKNOWN
G26 Registry upkeep	Who owns each approved tool, health checks, and retiring old versions.	Pam	Phase 3 (Dec)
G5, G7 Cost and ROI	How we track token spend and prove return; no numbers until the pilot measures them.	Oz; Pam; Sal	Phase 5 (from Dec)
G2, G6, G8, G13, G21, G27, G28 Later phases	Bot-to-bot handoffs, Cloud iQ signals, offline sites, noisy-neighbour limits, shared-hardware risk, unsanctioned AI use, and avoiding vendor lock-in.	Ed (mostly)	Phases 3-6 (Dec 2026 onward)

3. New connector gaps from today's governance check (G46-G52)

Check result: 0 MET, 5 PARTIAL, 1 MISSING (nothing reads from a real source yet). Pilot minimum: G46, G49, G50, G52.

Item	What it is / why it matters	Owner	Status
G46 Delegated-only hop	The connector reads only with the user's own sign-in, and results never exceed what the user can see (intersected with the SharePoint Sites.Selected grant).	Ed	Prototype built (connector-proto/, synthetic, unwitnessed); live test waits on SharePoint site + grant
G47 Source credentials	Per-client credentials for source systems, with expiry and revocation.	Ed	OPEN; waits on hosting/Key Vault
G48 Scope grants	Scopes granted per client and connector version; broad scopes rejected, writes off by default.	Ed	OPEN
G49 App-only limits	App-only access allowed only for metadata (health, grant status) and always flagged in the audit.	Ed	Prototype built
G50 Intake audit event	Every document gets an intake record before any model call; that record starts the audit clock and links through to the write.	Ed	Prototype built; live test waits on the shared intake mailbox
G51 Version pinning	Pin connector versions in the real gateway; check scope changes when a version bumps.	Ed	OPEN

Item	What it is / why it matters	Owner	Status
G52 Pilot sources	Which sources the pilot reads. Recommendation: SharePoint library via Sites.Selected + shared intake mailbox; QuickBooks excluded.	Jon	Decide Tue Oct 20

UNKNOWN = not in any file yet. Nothing here has been sent to anyone.