

Neatlogs Usability Report

How to read this

A fresh user signed up at `dev.neatlogs.com` on 27 May 2026 and walked the entire onboarding plus every empty state in the app. This report covers each step in order. For each step we give:

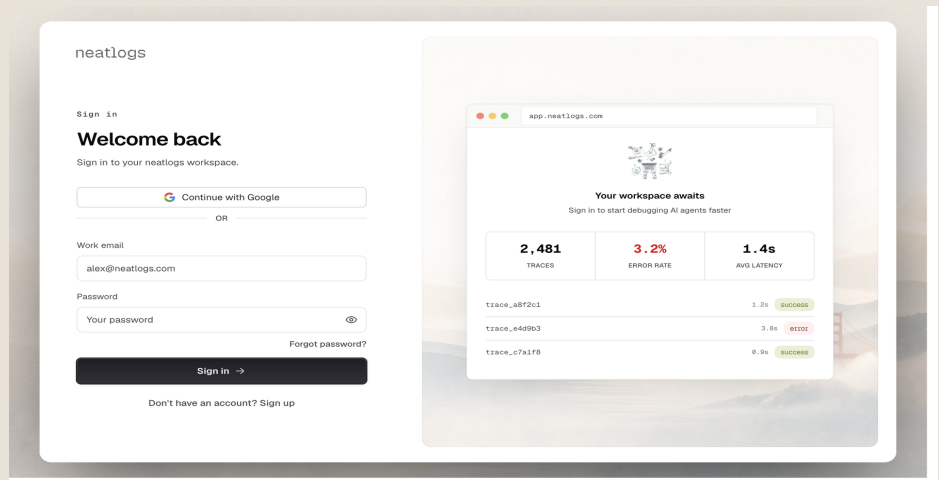
- **What we expected** — the mental model the user brought in.
- **What we actually found** — the observed behaviour, screenshots inline.
- **What a SOTA product does** — concrete examples from current best-in-class developer tools (Linear, Vercel, Stripe, Supabase, PostHog, Plaid, Notion, Resend, Clerk, LangSmith, Langfuse).
- **Severity & suggested fix** — where applicable.

Severity scale:

- **Blocker** — user is stuck, lost, or churns.
- **Major** — significant friction or confusion.
- **Minor** — paper cut.
- **Polish** — easy win, low pain.

Where a step had nothing notable, we've kept it to a sentence. Where a step has a real issue, we've gone deep.

Step 0 – Sign-in landing



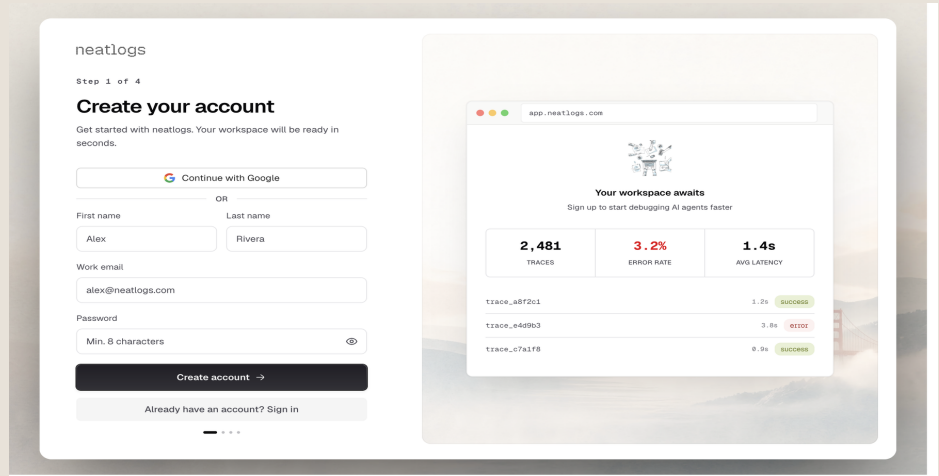
Expected. A standard sign-in screen with a clear path to sign-up if I'm new.

Actual. Got it. Email + password, Google OAuth, "Sign up" link visible. The right-side product preview (faux dashboard with metric tiles + trace rows) is well executed and sets expectations.

SOTA reference. Vercel, Resend, and Linear all use a similar split-screen with product imagery on the right; this matches the bar.

Severity. None – this step is good as-is.

Step 1 – Create your account



Expected. Name, email, password. Maybe a CAPTCHA or magic link.

Actual. First name, last name, work email, password. Eyebrow Step 1 of 4 – explicit funnel disclosure. Helpful.

Issues found:

1.1 No password requirements visible — Minor

The password field offers a show/hide eye toggle but no inline policy (“min 8 chars, must contain...”). We submitted Neat Test!2026 and it worked, but a user typing password123 won’t know if it’ll be rejected until they hit submit.

SOTA: Stripe, Linear, Clerk all show inline rules that turn green as each is satisfied. Supabase shows a strength meter. Even a simple “Min 8 characters, 1 number, 1 symbol” line below the field is enough.

Fix: add the policy as helper text below the password field, or render a live checklist as the user types.

1.2 “Work email” labelling but no enforcement — Minor

The label says “Work email” implying B2B, but accepts `colas.finned_1z@i-cloud.com` without warning. Either gate it, or use the label “Email”.

SOTA: Linear and Vercel say “Work email” and politely warn (not block) when a free-provider domain is used: “*We recommend using your work email so we can connect you with your team.*”

Fix: change the copy to “Email”, or add the soft warning. Don’t claim a constraint the system doesn’t enforce.

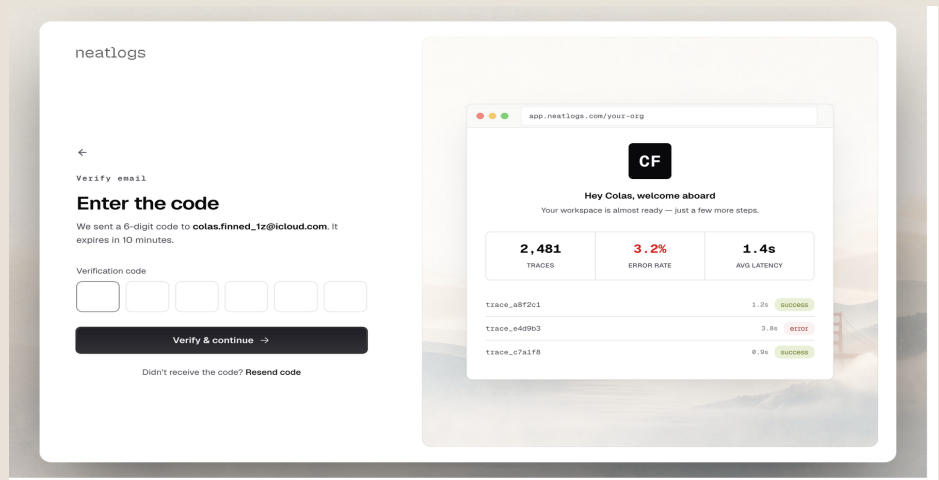
1.3 No visible stepper, only an eyebrow — Polish

ARIA snapshot shows `: navigation "Onboarding progress" landmark`, but at the captured viewport only the text “Step 1 of 4” rendered. Users want to see what’s ahead before committing to a multi-step flow.

SOTA: Plaid and Stripe Connect show a labelled segmented bar (Account → Verify → Workspace → Team). Even a thin progress bar at the top of the card is enough.

Fix: make the stepper visually present with step labels, not just a count.

Step 2 — Email verification



Expected. A 6-digit code, ~10-min expiry, paste-the-whole-code support, resend with cooldown.

Actual. Six individual digit inputs, 10-min TTL, resend button, helpful “expires in 10 minutes” copy. Standard.

Issues found:

2.1 Verify & continue not disabled when empty — Minor

The submit button accepts a click with all six digits empty (didn’t test the failure path explicitly, but no disabled state was visible during fill). It should be disabled until all six digits are entered.

SOTA: every OTP UI worth shipping (Stripe, Plaid, Clerk, Auth0) keeps the submit disabled until the input is valid.

Fix: disable until 6 digits entered.

2.2 No visible “paste full code” support — Minor

A real user usually copies the code from their email and pastes it. Six separate inputs without paste-distribution become a fight: you paste 355288 into the first box and end up with 3 in box 1, nothing in 2-6.

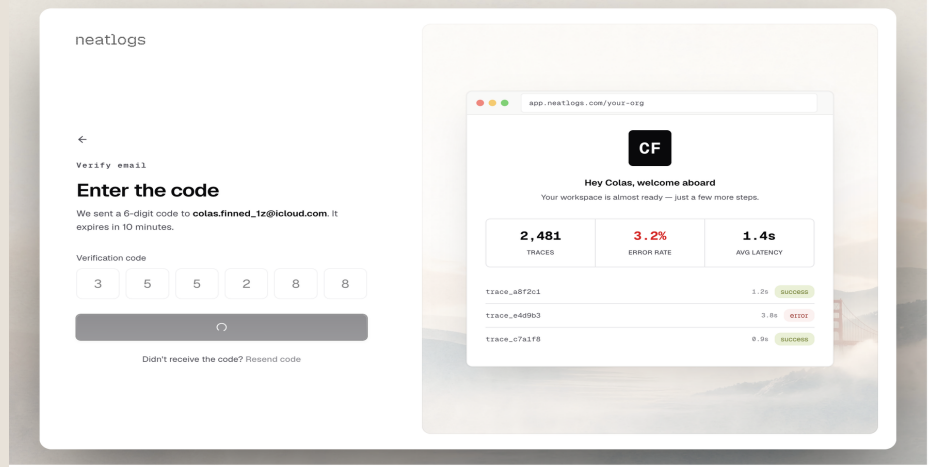
SOTA: Clerk, Linear, Vercel — pasting a 6-digit code into the first field auto-distributes across all six. This is now the expected behavior.

Fix: implement paste-distribution. (May already work — test and verify.)

2.3 No visible caps-lock or autofill hint — Polish

Not blocking, just a touch of polish. iCloud sometimes auto-fills the code from Mail; check that the inputs accept the autofill.

Step 3 — Workspace setup



Expected. “Name your team” or “Create your workspace” — one or two fields, get me to the product.

Actual. Three fields stacked: Organization name, Organisation URL, Workspace name. The eyebrow says Step 2 of 4. The slug auto-derives from org name (good).

Issues found:

3.1 Two entities created upfront with no explained difference — Major

The user is asked to name an **Organization** AND a **Workspace** on the same screen, with no explanation of what each is for. The helper text reads “*This is how your team will identify your organization in neatlogs*” — which only describes the org, leaving workspace dangling. After onboarding, the workspace name (Production) appears **nowhere** in the app — not in any breadcrumb, header, or setting we saw.

SOTA: Vercel asks for “Team name” only and creates a default project under it. Linear asks for “Workspace name” only. Supabase asks for “Organization name” then a separate “Project name” — but it ties them together visually and explains the relationship.

The user is being charged the cognitive cost of a multi-tenant data model that the SDK (per public docs) doesn’t even expose — the SDK groups by `workflow_name` keyed off API key, not by org/workspace.

Fix (recommended): drop the workspace step entirely. Org + Project is the model that the SDK and the app actually use. If multi-environment support is needed (prod / staging / dev), express it as a Project property or as a workspace switcher in the topbar — not as an upfront question.

Fix (minimum): if both must stay, add a helper sentence under each: “*Organization is your billing entity. Workspace is an environment within it (e.g., production).*” And **show the workspace name** somewhere in the post-onboarding UI.

3.2 Organization name and Organisation URL — locale mismatch — Polish

Organization (en-US) sits directly abc Organisation URL (en-GB). Reads as careless.

Fix: pick one. en-US recommended given the app’s other copy.

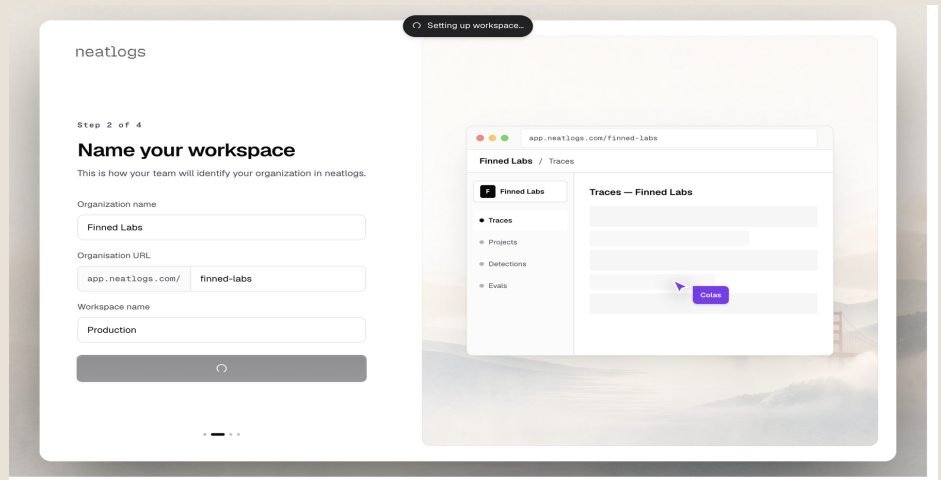
3.3 URL slug field has no live availability check — Minor

The slug `finned-labs` was accepted; we don’t know if a duplicate would show inline error or fail at submit. SOTA tools show a green check next to a free slug and a red `⊗` `already taken` for collisions, debounced.

SOTA: Vercel, GitHub, Linear, Notion all do live slug availability checks.

Fix: add debounced async availability check.

Step 4 — “Setting up workspace...”

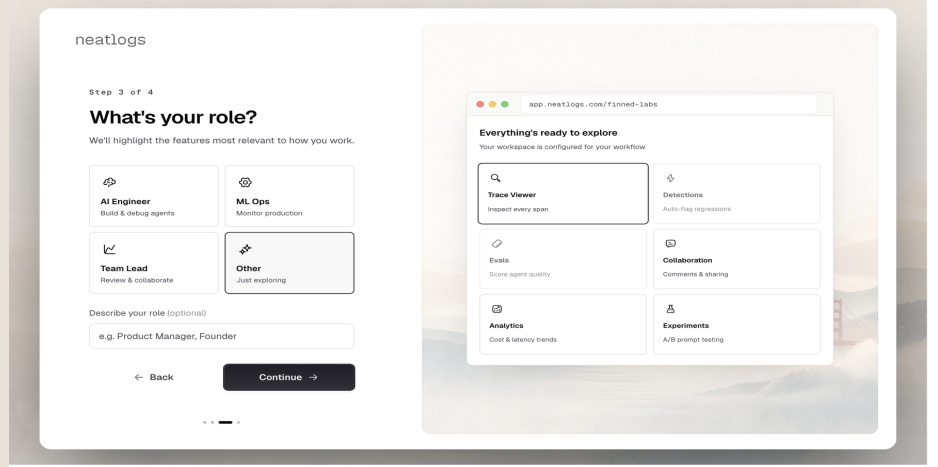


Expected. A loading state while the org is provisioned.

Actual. Got it — Continue button becomes a spinner, separate “Setting up workspace...” message renders. The right-side preview updates to reflect the typed values *before* the API call returns. That’s a nice touch.

Severity. None — this step is good.

Step 5 — Role qualifier



Expected. “Tell us how you’ll use this so we can tailor the experience” — and then the experience actually gets tailored.

Actual. Four role cards (AI Engineer / ML Ops / Team Lead / Other), optional free-text. We picked AI Engineer.

Issues found:

5.1 The role we picked has no visible effect — Major

After picking AI Engineer, nothing in the post-onboarding app appeared role-specific. Same nav, same empty states, same default Settings tab, same suggested prompts. If a user is asked to declare intent, the product owes them at least one tailored thing in return.

SOTA:

- **Notion** asks if you’re using it for personal/team/work and changes the first-run templates.
- **Linear** asks team size and pre-creates appropriate views.
- **Stripe** asks business type and wires the right tax/compliance defaults.
- **PostHog** asks role and changes the dashboard’s first cards.

All of them deliver on the question they asked.

Fix: use the answer. Concrete suggestions:

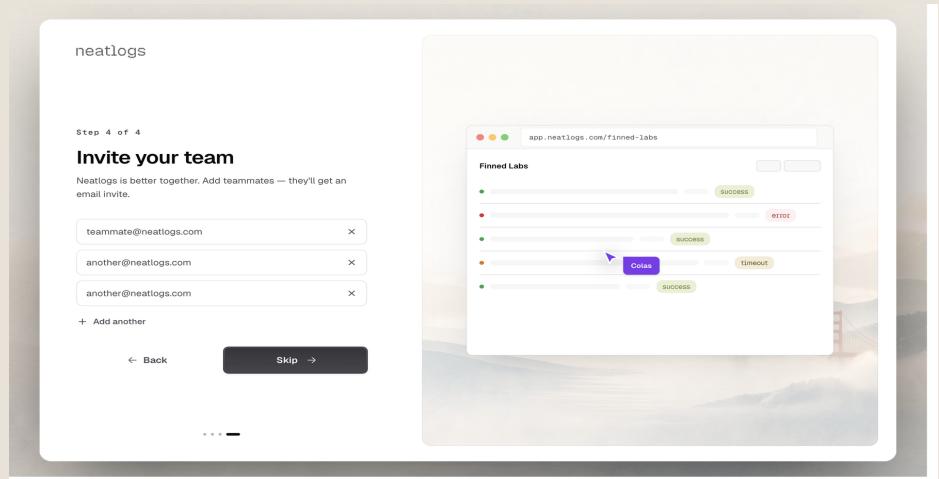
- AI Engineer → land them on a “Connect your first agent” screen with a Python SDK snippet.
- ML Ops → land them on Analytics with a “We’ll show production metrics here” empty state.
- Team Lead → land them on Triage with an explanation of the workflow.
- Other → land them on Investigate with educational content rather than data-assuming prompts.

5.2 Free-text “Describe your role” with no apparent use — Minor

If the four canonical roles are enough, drop the free-text. If they aren't, do something with the input (route to onboarding emails, surface in Activity, etc.). Right now it looks like a survey field that goes into a void.

Fix: remove, or use it.

Step 6 — Invite team



Expected. Optional team invites, easy to skip, easy to come back to.

Actual. Three pre-rendered email rows, an “Add another” button, and a primary CTA labelled Skip.

Issues found:

6.1 The primary CTA is “Skip” — Blocker (for activation, not for the user)

The most prominent button is *don't do this step*. That's a self-aware admission that the step is in the wrong place. If most users will skip, why is it step 4 of 4 — i.e., the last thing between them and the product?

This is the core activation issue. The user finished onboarding without the one thing that unlocks the product: **a project**. Six of eight nav items in the post-onboarding app render No project selected. The most important first action is buried in the sidebar as a list-item-shaped button.

SOTA:

- **Vercel** auto-creates a project during signup if you connect a Git provider; otherwise the post-signup screen IS the new-project screen.
- **PostHog** ends onboarding with the SDK install snippet on screen, key already generated, waiting for first event.
- **Resend** ends onboarding with Send your first email — the product's primary verb is the next thing you do.

- **Linear** lands you in a project with sample issues so the empty state is never empty for a new user.
- **Supabase** auto-creates a default project so you immediately have a database connection string.

The pattern is consistent: the activation step (first project / first event / first send) belongs **inside** onboarding, not after it.

Fix:

1. **Auto-create a default project** at the end of step 3, named after the workspace or org slug. Show its name and let the user rename inline.
2. **Replace step 4 with: “Connect your first project.”** Show the API key and a copy-pasteable SDK snippet for the role they picked in step 5. Show a live “waiting for first trace” indicator.
3. **Demote** `Invite team` to a dismissible card alongside the SDK snippet, or to the empty state on the (now project-aware) Triage / Analytics page.

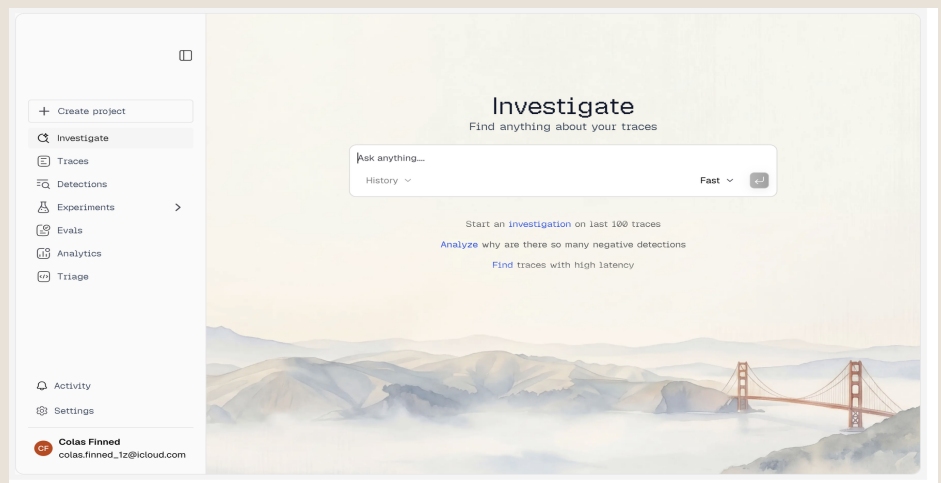
6.2 No way to bulk-paste emails — Polish

Three pre-filled rows + one “Add another” button. A user with a list of 10 teammates is stuck clicking ten times.

SOTA: Linear, Notion, Vercel all accept comma- or newline-separated emails in a single field.

Fix: allow paste-of-list, parse on the fly.

Step 7 — Post-onboarding landing (Investigate)



Expected. A welcome screen, a quickstart card, or a “create your first project” prompt. Something pointing me at first value.

Actual. Landed at /, which renders **Investigate** — an AI search/chat interface. Three suggested prompts:

- “Analyze why are there so many negative detections”
- “Find traces with high latency”
- “Summarize user checkout flows”

A fresh user has zero traces, zero detections, and no checkout flow. Every suggestion is a dead end.

Issues found:

7.1 Landing page assumes data the user doesn’t have — Blocker

This is the moment a brand-new user forms an opinion of the product. They land on a search box that searches nothing, with prompts that invite questions about data that doesn’t exist, with a sidebar full of items that all say “No project selected.” The product looks broken on first impression.

SOTA:

- **PostHog:** post-onboarding lands on “Install snippet” with the install code, key, and a live “waiting for first event” indicator. Once the event arrives, the page transitions to the dashboard.
- **Sentry:** post-onboarding lands on the project’s installation screen with platform picker, copy-pasteable snippet, and `Waiting for first error` widget. Once the error fires, the page swaps to issues.
- **LangSmith / Langfuse:** post-onboarding shows the API key + SDK snippet up front. Investigate-style features come later.

Notice the pattern: the landing page is a **setup page until first data**, then it becomes the analytics page. State-aware routing.

Fix: route logged-in users with zero data to a dedicated `/setup (c /projects/new)` screen. Once the first trace arrives, route them to Investigate or Traces. Investigate as a feature is fine; as the empty-state landing it’s wrong.

7.2 Suggested prompts are static — Major

The three suggestions don’t adapt to data presence. Even after data exists, they’ll likely be generic. SOTA AI search tools (Linear’s AI search, Notion AI, GitHub Copilot Chat) tailor suggestions based on what you actually have.

Fix: state-aware suggestions. Zero data → “How do I send my first trace?” / “Show me the SDK quickstart” / “What does Neatlogs do?”. With data → tailored to recent traces, recent detections, recent latency outliers.

7.3 Create project is a sidebar list-item — Major

The single most important action in the app is rendered with the same visual weight as ordinary nav items.

SOTA: PostHog, Sentry, Linear all promote the “next setup step” with a distinct visual treatment — a banner, a callout card, or a button with high contrast.

Fix: when zero projects exist, replace the `Create project` list item with a prominent button (or a topbar banner). When `1` project exists, demote it.

Step 8 — Empty states across the nav

Expected. Either a “no data yet” state with a clear next action, or content if there’s any.

Actual. Six routes (`/traces`, `/detections`, `/experiments`, `/evals`, `/analytics`, `/triage`) all render the same illustration with the headline `No project selected` and helper text `Select a project from the sidebar to view traces`.

Issues found:

8.1 Empty-state copy points to a thing that doesn’t exist — Major

“*Select a project from the sidebar*” implies there is a list of projects to select from. There isn’t. The user reads this six times across six routes and gets six dead ends.

SOTA:

- **Sentry** empty state when no project: `Create your first project` with a button.
- **Vercel** empty state when no project: `Deploy your first project` with three quickstart paths.
- **Linear** empty state in any view with no data: shows a contextual CTA (“`Create your first issue`”) with a sample-data toggle.

Fix: state-aware empty state. When zero projects exist, the headline becomes `Create your first project to see traces` with an inline CTA. When `1` project exists, the current copy is fine — but the action `Select a project from the sidebar` only makes sense if there *is* a project list in the sidebar. Today there isn’t, so this copy never makes sense.

8.2 Same dead-end repeated six times — Major

The empty state component is shared (good — single fix). But every route shows the same message. There's no per-route educational content, no "here's what Detections will look like once you have data" preview.

SOTA: Linear, Notion, Stripe all use the empty state as a teaching moment — illustrations of what the page will look like, sample data toggle, or video walkthrough.

Fix: per-route empty-state copy and (optional) a sample-data toggle that swaps in synthetic traces / detections / etc. so the user can explore the product before instrumenting.

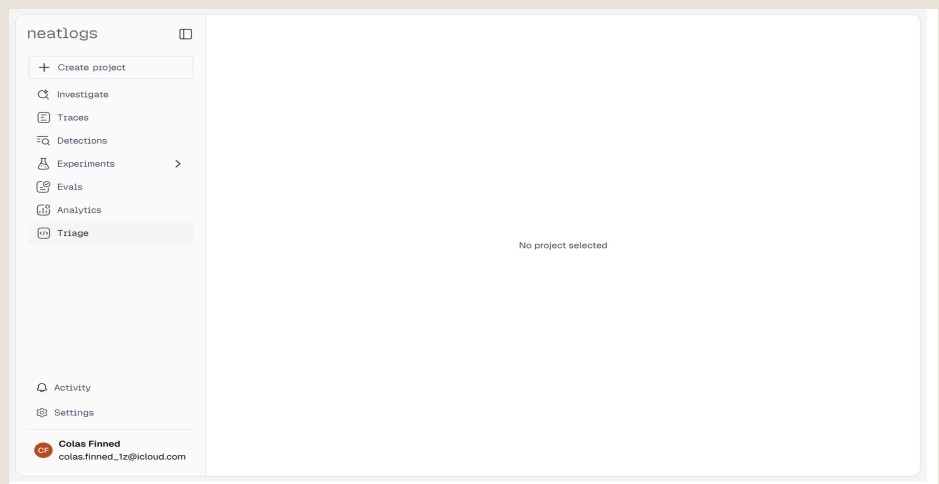
8.3 No in-app help / docs / quickstart link — Major

Across all six dead-end routes, there is no link to docs, no ? help button in the topbar, no quickstart card. The user's only escape is to leave the app and Google.

SOTA: every developer tool worth its salt has a persistent ? button or Help slide-over (Stripe, Linear, Vercel, Supabase, PostHog).

Fix: add a persistent ? in the topbar (or user chip) that opens the Quickstart in a slideover. Cost: trivial. Benefit: huge.

Step 9 — Activity (sidebar item)



Expected. Audit log? Live event feed? Notification center? Some kind of activity stream.

Actual. Clicking it does nothing visible. No popover, no panel, no expansion in the snapshot. Either it's broken, hidden behind data, or simply unimplemented.

Issues found:

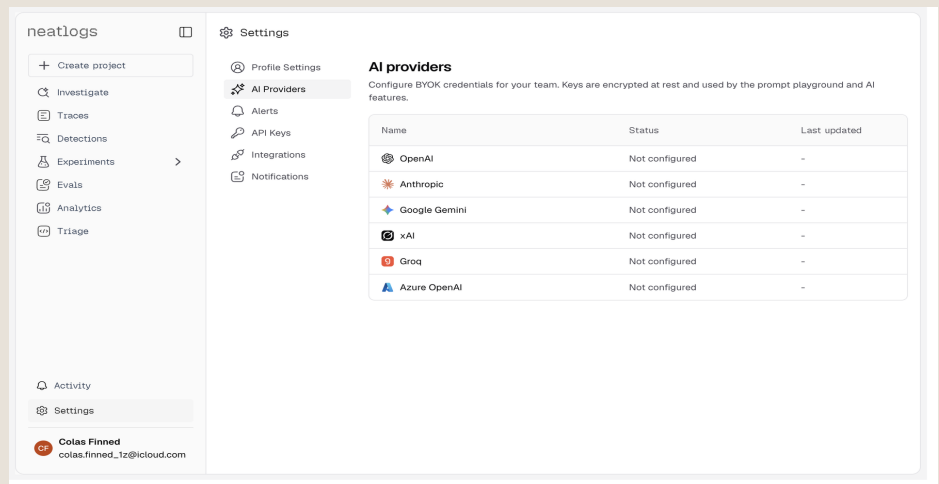
9.1 A nav item that doesn't respond — Major

For a brand-new user, an unresponsive nav item reads as broken. This is the kind of thing that lowers trust in the *rest* of the product.

SOTA: Linear, Notion, Vercel — every nav item either does something visible immediately or has a clear empty state. None silently no-op.

Fix: either implement it, hide it until it has content, or render a “Nothing here yet — Activity will show audit events as your team works” empty state.

Step 10 — Settings → AI Providers (default tab)



Expected. Profile or General as the default Settings tab.

Actual. **AI Providers** is the default tab. Six BYOK rows (OpenAI, Anthropic, Google, xAI, Groq, Azure OpenAI), all Not configured.

Issues found:

10.1 BYOK setup is buried in Settings, but is genuinely needed for AI features — Major

The helper copy on this page says “*Keys are encrypted at rest and used by the prompt playground and AI features.*” So Investigate (the AI search landing page) likely needs at least one BYOK to function. But onboarding never mentioned this. We just spent five minutes onboarding and the AI features can't actually work yet because we haven't given a key.

SOTA:

- **Cursor** asks for an API key (or uses their hosted key) during first-run.
- **Continue.dev** asks during install.
- **PostHog** asks for OpenAI key when you first try to use AI features, with a clear modal.

Fix: either:

- **Promote BYOK setup into onboarding** (after role qualifier, before invite team — make it step 4, optional with a default fallback).
- **Or detect missing BYOK** when the user first tries Investigate and show an inline “configure a provider to use this” prompt with a one-click jump to settings.

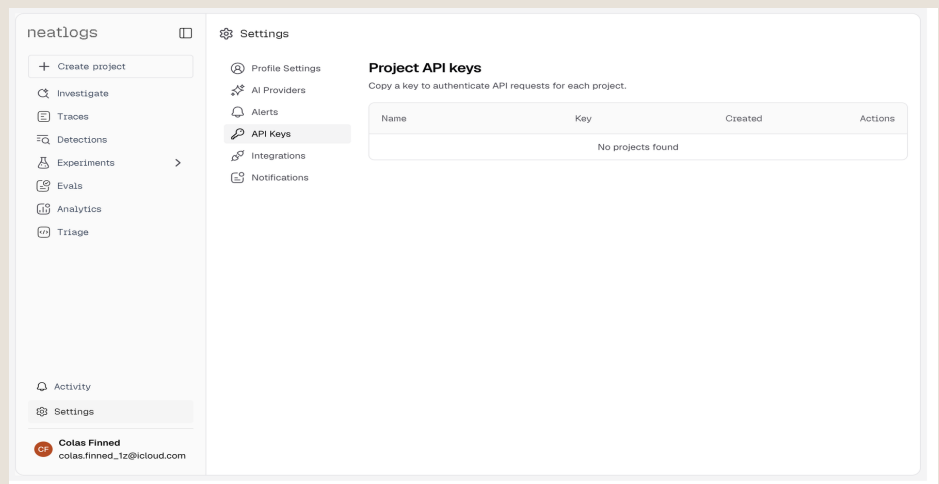
Today, the user can’t tell why Investigate isn’t useful. They just see suggestions that go nowhere.

10.2 Six providers shown with no recommendation — Polish

If the user is new and just wants the AI features to work, they don’t know which provider to pick. SOTA tools either default-recommend one (“Most users start with OpenAI”) or let the user use a hosted key with a small free tier.

Fix: mark a recommended provider, or offer a hosted-key default with quotas.

Step 11 — Settings → API Keys



Expected. A list of project API keys with copy buttons.

Actual. Empty table with the cell `No projects found`.

Issues found:

11.1 The page that gates SDK integration says “No projects found” — Blocker

This is the second-most-important page in the product (after the dashboard) — it’s where users grab the credential they need to send their first trace. And it’s a dead end because no project exists. There’s no inline `Create project` CTA on this page; the user has to back out and find the sidebar button.

SOTA: Stripe's API keys page never shows an empty state — keys exist by default for the test mode. PostHog's project settings page has the install snippet directly. Sentry pre-creates a default project.

Fix: when zero projects exist, replace the “No projects found” cell with an inline `Create your first project` CTA that opens the project-creation flow modally and returns the user to this page with the new key revealed.

Step 12 — Triage vs Detections (and Evals vs Experiments)

These are sidebar siblings. From the outside, the boundaries are unclear.

Issues found:

12.1 Two sets of overlapping concepts in primary nav — Major

- **Detections** (likely: automated rule findings).
- **Triage** (likely: queue of things needing human action).

If Triage is a queue *over* Detections (and over comments / threads / open issues), it's a workflow surface, not a peer entity. Putting them side-by-side in the nav implies they're equivalent things.

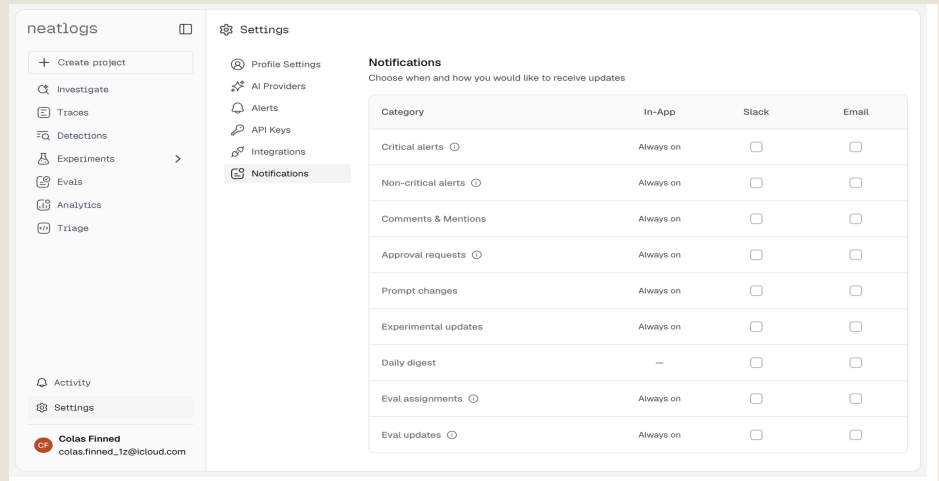
Same problem with **Evals** (scoring rubrics) and **Experiments** (A/B prompt variants scored by Evals). Experiments uses Evals; they're not peers.

SOTA:

- **Sentry** has Issues (the data) and Inbox (the queue) — but it's named **Inbox**, which makes the workflow nature obvious.
- **GitHub** has Issues (data) and Notifications (queue) — same pattern, clearer naming.
- **Linear** consolidates queues into a single Inbox per user.

Fix: rename Triage → Inbox or Queue, and consider nesting it visually under Detections. For Evals/Experiments, either nest Experiments under Eval or rename to clarify the relationship.

Step 13 — Settings → Notifications



Expected. A matrix of notification types vs channels.

Actual. Got it. `Critical alerts` is locked to “Always on” with an info icon. Other categories presumably toggleable.

Severity. None notable — the matrix is well structured. Polish ask: ensure the locked row’s “Always on” cells visually communicate “this can’t be changed” rather than just being plain text (a small lock icon would help).

Cross-cutting issues

Issues that span multiple screens.

C.1 Workspace name asked, never used — Major

The user typed `Production` in step 3. This name appears nowhere in the post-onboarding app. Either the workspace concept is real and should be visible (topbar switcher, breadcrumbs, settings), or it’s not real and shouldn’t be asked for.


C.2 Terminology drift between SDK and UI — Major

Per public docs, the SDK uses `workflow_name` to group traces. The UI calls this concept “Project”. A user moving from UI to SDK on day two will hit “wait, where’s `project_name`?” friction.

SOTA: Stripe, Vercel, Supabase — terminology is identical across UI, SDK, and docs. When they migrate names, they alias the old one for a long deprecation window.

Fix: pick one term. “Project” is the right choice for users. Add a `project_-name` alias in the SDK; deprecate `workflow_name` in docs over time.

C.3 No visible relationship between SDK install and the app — Major

The app never tells the user “to send your first trace, do X”. There’s no SDK snippet anywhere visible in the post-onboarding flow,  Connect your app step, no install widget on any empty state.

SOTA: PostHog, Sentry, Resend, Stripe all ship the install snippet *visibly inside the app*, prominently, with the user’s actual key pre-filled. Copy-paste, run, see the result.

Fix: add an “Install SDK” surface to the empty state and to the project page, with the user’s API key pre-filled in a Python (and eventually JS) snippet.

C.4 No demo / sample-data option — Polish

A user evaluating the product can’t see what Detections or Analytics look like populated without instrumenting their own app first. That’s a high cost-to-evaluate.

SOTA: Linear has a “Try with sample data” toggle. Vercel has demo deployments. Sentry has a demo project on every account. Notion has template galleries.

Fix: offer a Demo workspace toggle that fills in synthetic traces / detections / evals so users can navigate the populated product before committing to instrumentation.

C.5 No status, doc, or help affordance in-app — Major

Already covered in 8.3 — no ?, no status link, no doc slideover. Adding one would close many of the smaller gaps in this report at low cost.

Severity summary

| Step / area | Issue | Severity |
|-------------|--|--------------------------------|
| 6.1 | Onboarding ends at skip instead of project setup | Blocker (activation) |
| 7.1 | Investigate as zero-data landing | Blocker (impression) |
| 11.1 | API Keys page empty with no inline CTA | Blocker (path to value) |
| 3.1 | Workspace asked but never used / explained | Major |
| 5.1 | Role qualifier has no effect | Major |
| 7.3 | Create project is a list-item, not a CTA | Major |

| Step / area | Issue | Severity |
|-------------|---|----------|
| 8.1 | Empty-state copy refers to non-existent project list | Major |
| 8.2 | Six routes share one dead-end | Major |
| 8.3 | No in-app docs / help | Major |
| 9.1 | Activity does nothing | Major |
| 10.1 | BYOK needed for AI but not surfaced in onboarding | Major |
| 12.1 | Triage/Detections and Evals/Experiments seams unclear | Major |
| C.1 | Workspace name vanishes | Major |
| C.2 | SDK workflow vs UI Project terminology drift | Major |
| C.3 | SDK install never surfaced in the app | Major |
| C.5 | No ? help affordance | Major |
| 1.1 | Password policy not visible | Minor |
| 1.2 | “Work email” labelling, no enforcement | Minor |
| 2.1 | OTP submit not disabled when empty | Minor |
| 2.2 | OTP paste-distribution unclear | Minor |
| 3.3 | Slug uniqueness not checked live | Minor |
| 5.2 | Free-text “Describe your role” with no use | Minor |
| 6.2 | No bulk-paste in invite list | Polish |
| 1.3 | Stepper invisible at viewport | Polish |
| 3.2 | en-US / en-GB locale mix | Polish |
| 10.2 | No recommended BYOK provider | Polish |
| 13 | Locked-row “Always on” needs lock icon | Polish |

| Step / area | Issue | Severity |
|-------------|------------------------------|----------|
| C.4 | No demo / sample-data option | Polish |

What we'd ship first

If we could only do five things this quarter to fix the activation gap:

1. **Auto-create a default project at end of step 3.** Replace step 4's "Invite team" with "Connect your first project" — show API key, SDK snippet, and a live "waiting for first trace" indicator. (Fixes 6.1, 7.1, 11.1, C.3.)
2. **State-aware empty states.** When zero projects exist, every empty state becomes "Create your first project" with an inline CTA. When 1 project exists, copy stays as-is. (Fixes 8.1, 8.2.)
3. **Persistent in-app help.** Add ? in topbar opening Quickstart in a slideover. (Fixes 8.3, C.5.)
4. **Surface or kill Workspace.** Either add a workspace switcher to the topbar, or remove the field from onboarding. Reconcile SD workflow_name ↔ UI Project. (Fixes C.1, C.2, 3.1.)
5. **Make role pick matter.** Use the answer to pick the SDK snippet, the landing route, and the default Settings tab. (Fixes 5.1.)

Everything else in this report is real but secondary. Fix these five and the product will *feel* like it shipped a major release, even though most of the change is rearrangement and copy.