AmberTrace

User Guide

Building explainable, verified AI over your own expertise — for domain experts, not engineers.

Product: AmberTrace — app.ambertrace.ai

Audience: Domain experts (clinicians, risk & compliance leads, analysts)

Covers: Every screen you see, Query vs Predictions, and permissioning an AI agent

From: Ambertrace Labs Ltd. — UK Sovereign AI

AMBERTRACE LABS

AmberTrace — Legal Notices

User Guide · Ambertrace Labs Ltd · August 2026

Please read before use. AmberTrace is a professional decision-support tool. The notices below apply to all use of the platform, including clinical, financial, regulatory, and autonomous-agent workflows. By using AmberTrace you agree to these terms and to the Master Services Agreement applicable to your account.

Not Professional Advice

AmberTrace outputs — including answers, forecasts, rules, confidence scores, verified results, and proof certificates — are AI-assisted decision support only. They do not constitute medical advice, clinical diagnosis, legal advice, financial advice, or investment advice. No output should be acted upon without appropriate human review by a qualified professional in the relevant field. Users and their organisations remain solely responsible for all decisions made using AmberTrace.

Regulatory Status

AmberTrace is not a regulated medical device and has not been assessed or cleared by the MHRA, the FDA, or any other health-technology regulatory authority. It must not be used as a substitute for clinical judgment or as a standalone decision-making system in any context where regulatory clearance as a medical device, in-vitro diagnostic, or clinical decision-support tool would be required.

AmberTrace is also not authorised or regulated by the Financial Conduct Authority (FCA), the U.S. Securities and Exchange Commission (SEC), or any other financial regulatory authority in any jurisdiction. The platform is an AI software tool and is not a regulated financial service, a financial instrument, or investment, trading, or financial advice.

Organisations deploying AmberTrace for regulated activities — including clinical triage, loan decisioning, insurance underwriting, or investment and trading decisions — remain responsible for ensuring their own compliance with all applicable regulatory frameworks (including the UK AI Act, EU AI Act, MHRA guidance on Software as a Medical Device, and, for financial use, FCA/PRA, MiFID II model-risk and SEC requirements), and for the governance and validation of any model used in regulated decision-making.

Accuracy and Scope of “Proof”

The terms “verified” and “proof” as used in this guide refer to formal verification that an answer provably follows the rules you have written. They do not guarantee factual accuracy, clinical correctness, or fitness for any particular purpose. AmberTrace’s outputs are probabilistic and depend on the quality of the data and rules you supply. A verified answer with high confidence is not a guarantee of a correct outcome.

Data Protection

When you upload data to AmberTrace, your organisation is the data controller. Ambertrace Labs Ltd acts as a data processor and processes your data only in accordance with a Data Processing Agreement (DPA), available on request. You are responsible for ensuring you have a valid lawful basis under UK GDPR before uploading data relating to identifiable individuals. Health and financial data carry additional legal obligations — please consult your Data Protection Officer or legal counsel before uploading such data. Data-protection enquiries: legal@ambertracelabs.com.

Intellectual Property

Rules, platforms, and models you create from your own data are yours. Ambertrace Labs Ltd retains all intellectual property rights in the AmberTrace platform, its proprietary models, and the methods used to generate rules and forecasts. For full IP terms, see your Master Services Agreement.

Confidentiality & Permitted Use

This guide and the AmberTrace platform contain material proprietary to Ambertrace Labs Ltd. This guide is provided for the use of authorised users of the platform and may not be reproduced, distributed, or disclosed to any third party without the prior written consent of Ambertrace Labs Ltd, except as permitted under your Master Services Agreement.

Governing Law

AmberTrace is operated by Ambertrace Labs Ltd, registered in England and Wales (Company No. 15396088). Use of this platform is governed by English law. The Master Services Agreement, Data Processing Agreement, and Privacy Policy are available at ambertrace.ai/legal or from your account manager.

Copyright

What's inside

GETTING STARTED

1. Welcome & the big idea

AmberTrace turns *your* expertise into an AI you can actually trust — one that explains every answer and never hides its reasoning.

Most AI tools give you an answer and ask you to take it on faith. AmberTrace is different. It combines two things:

- **Patterns learned from your data** — the "neural" part, which spots trends and signals.
- **Rules you define from your knowledge** — the "symbolic" part, which encodes how decisions *should* be made in your field.

Putting the two together is called **neurosymbolic AI**. The pay-off: every answer comes with a plain-English **explanation** of how it was reached, and — where it matters most — a mathematical **proof** that the answer obeys your rules.

> **THE JOURNEY IN FOUR STEPS**
>
> You'll move through four simple stages, and the app guides you the whole way:
>
> **1. Domain** (describe your field) → **2. Data** (upload your information) → **3. Platform** (build the AI) → **4. Console** (ask it questions). Each screen even shows you which step you're on.

You do not need to write any code. You describe what you know in plain language, upload a spreadsheet, and AmberTrace assembles the rest. This guide walks you through every screen you'll see, then explains the two headline tools — **Query** and **Predictions** — and how to let an AI assistant use your platform on your behalf.

> **THE MODEL BEHIND IT — AND WHERE IT RUNS**
>
> The neural side of AmberTrace is powered by our own **proprietary model** — built on the open **Llama** family and then significantly enhanced by Ambertrace Labs for this work. It runs entirely on **Ambertrace's own UK-sovereign infrastructure**: your data is never sent to a third-party AI provider, and is never used to train anyone else's model. (More on security in Section 18.)

> **CAPABILITY GATING — SOME SCREENS DEPEND ON YOUR ORGANISATION'S PLAN**
>
> Your organisation's administrator controls which capabilities are enabled. Some screens — **Chat**, **Console**, and **Predictions** — only appear in the navigation when they are switched on for your organisation. If you do not see one of these screens, it is because your organisation's capability settings do not currently include it. Speak to your organisation administrator if you need access.

GETTING STARTED

2. Signing in

Everything lives at `app.ambertrace.ai`. You sign in with your username (or email) and password, or via your organisation's Single Sign-On (SSO) provider.

1. Open `https://app.ambertrace.ai` in your web browser.
2. Enter your **username or email** and your **password**. Click *Show* if you want to check what you typed.
3. Click **Sign In**. You land on your workspace, starting at the Chat screen (if enabled for your organisation).

If your organisation uses **Single Sign-On**, click **Sign in with SSO** instead and follow your identity provider's flow. If your administrator has enabled **multi-factor authentication** (MFA), you will also be prompted for a one-time code from your authenticator app.

Forgotten your password? Use the **Forgot password?** link to receive a reset email. New to AmberTrace? **Create an account** starts the sign-up flow.

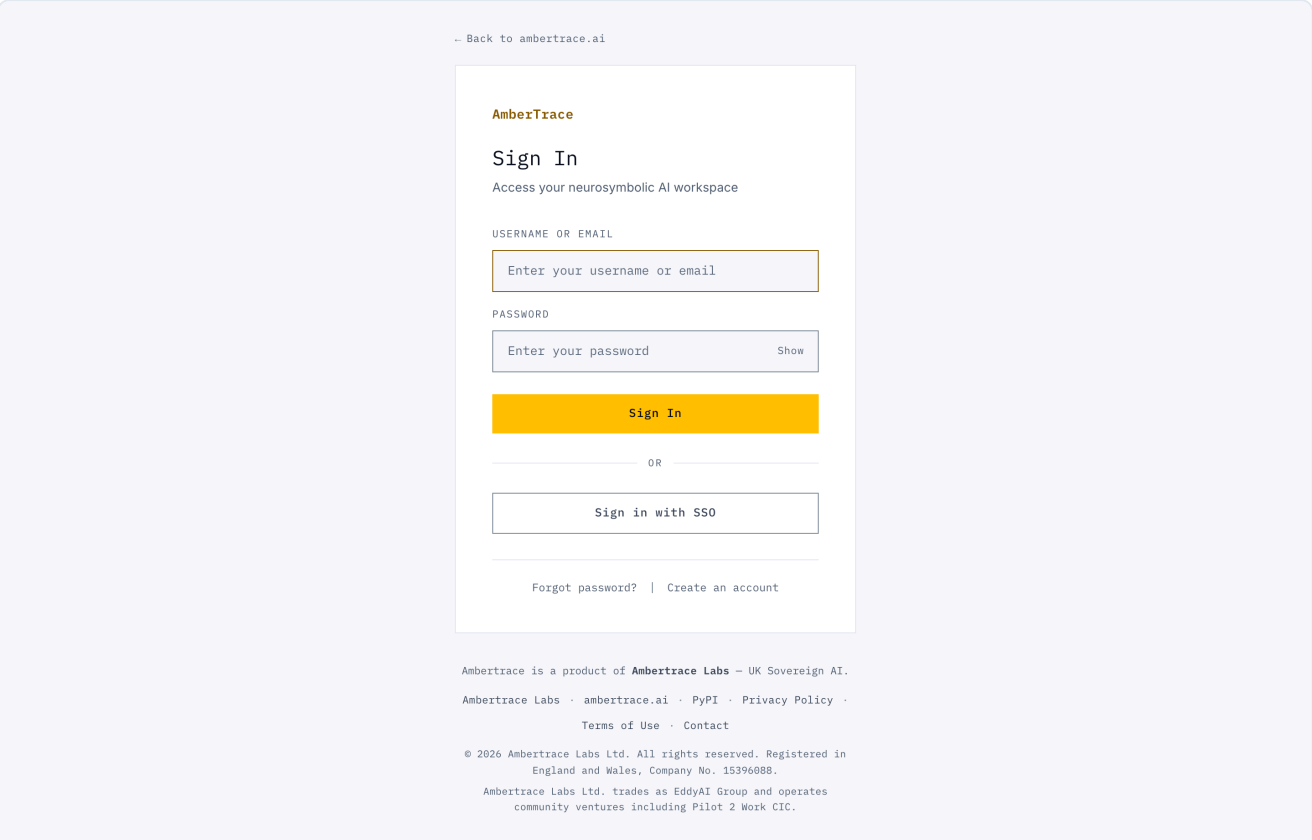

The sign-in screen at app.ambertrace.ai, with password and SSO options.

Once you're in, a slim icon bar runs across the top of every page. Hovering over an icon shows its name. Those icons are the screens this guide covers next. Your name sits on the right — click it for **Settings**, or use **Sign Out** to leave. A light/dark theme toggle sits beside your name.

YOUR SCREENS, ONE BY ONE · CAPABILITY-DEPENDENT

3. Chat

Your guided assistant. The fastest way to get started — just describe what you want to build, in your own words.

What it's for

Chat is a conversation with AmberTrace's assistant. Instead of clicking through every screen yourself, you can tell it what you're trying to achieve and it will help you set things up — from describing your field through to building a working platform.

> **CAPABILITY-DEPENDENT**
>
> Chat appears in the navigation only when your organisation's **chat** capability is enabled. If you do not see it, your organisation administrator has not yet switched it on.

When you'd use it

Use Chat when you're starting out and aren't sure where to click, or when you'd simply rather explain your goal and let the assistant do the legwork.

Key actions

1. Type what you want to build in the box at the bottom — for example, *"I need an AI that triages referrals by severity."*
2. Or click one of the suggested starter prompts to see it in action.
3. Read the assistant's reply. It may carry out tasks for you (you'll see a small "Working…" indicator while it does) and explain what it set up.
4. Past conversations appear in the left-hand sidebar so you can pick up where you left off.

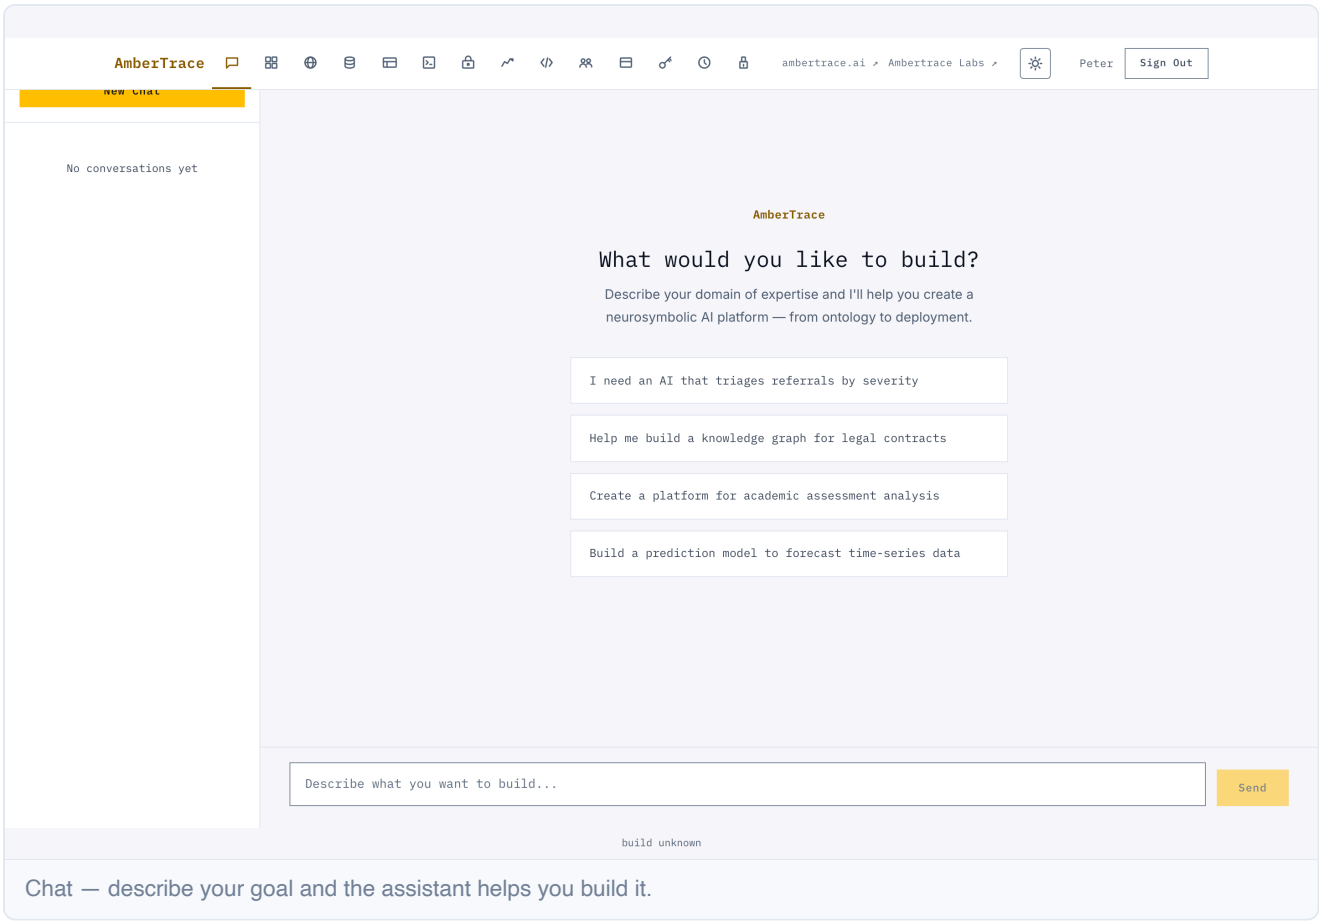

Chat — describe your goal and the assistant helps you build it.

YOUR SCREENS, ONE BY ONE

4. Overview

Your dashboard — a single glance at everything you've created and what to do next.

What it's for

The Overview (the dashboard) is your home base. It summarises your domains, datasets, and platforms, and points you toward the next sensible step.

When you'd use it

Drop in here to check the state of your work, see how many platforms you have running, or find the quickest route to your next task.

Key actions

- Review your counts of domains, datasets, and platforms at a glance.
- Follow the suggested next step (for example, "upload data" or "build a platform").
- Jump straight to any screen from here.

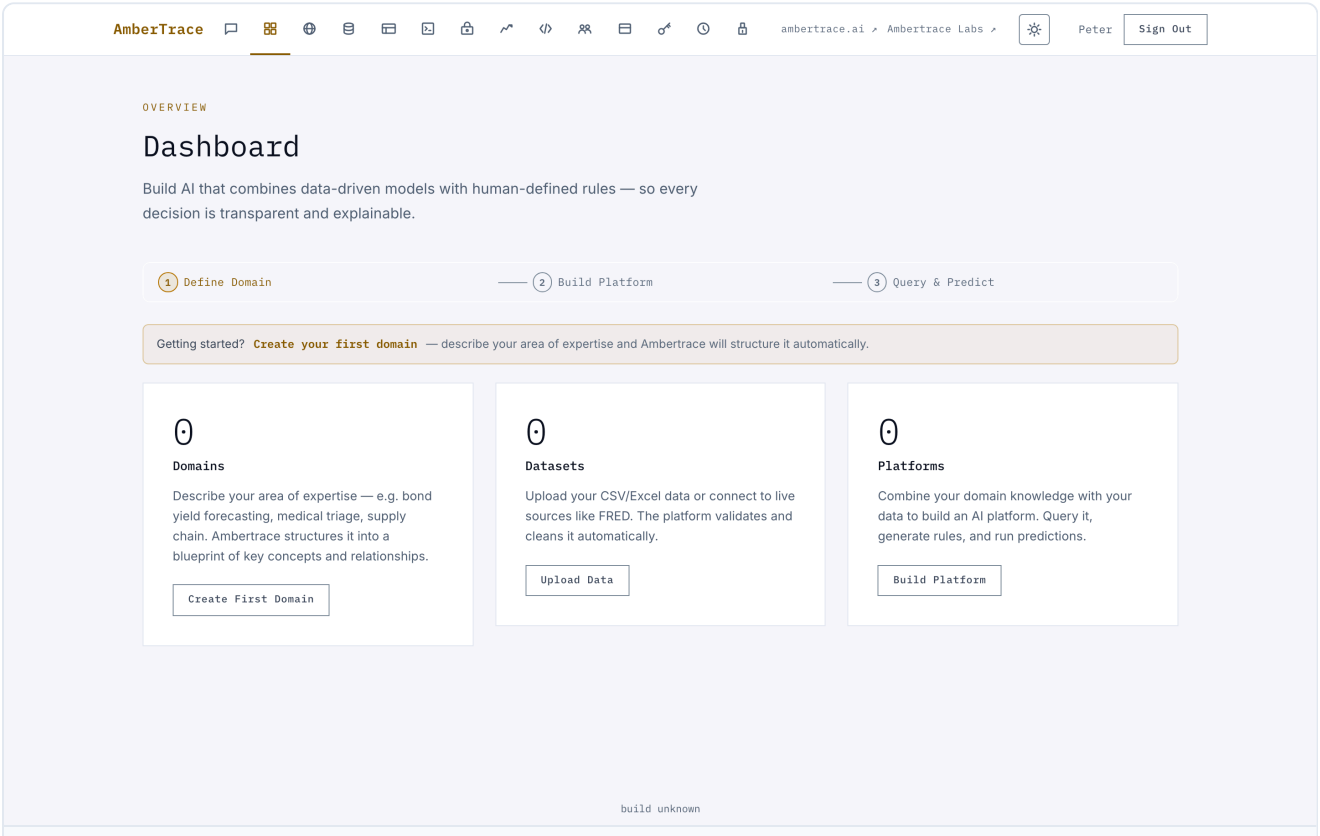

Overview — combine data-driven models with your own rules, transparently.

YOUR SCREENS, ONE BY ONE · STEP 1 OF 4

5. Domains

Describe what you know best. A “domain” is simply your area of expertise — the field, the key decisions, and the factors that matter.

What it’s for

This is where you tell AmberTrace about your world. You describe your field of expertise and the decisions you make, and AmberTrace structures that into a **blueprint** (technically called an *ontology*) that will power your AI.

When you’d use it

This is your **first step**. Create one domain per area of work — for example “Medical Triage” or “Legal Contract Review.”

Key actions

1. Give your domain a **name** (e.g. “Medical Triage”).
2. Write a **description** — the field, the key concepts, and what decisions need to be made. The more you say, the better the blueprint.
3. Click the amber button to create it. AmberTrace builds the blueprint for you.
4. Your domains appear in a list below, where you can **Edit**, review, or **Delete** them.

> **TIP**
> Write the description the way you’d brief a new colleague: what you assess, what information you look at, and how you decide. Plain language is perfect.

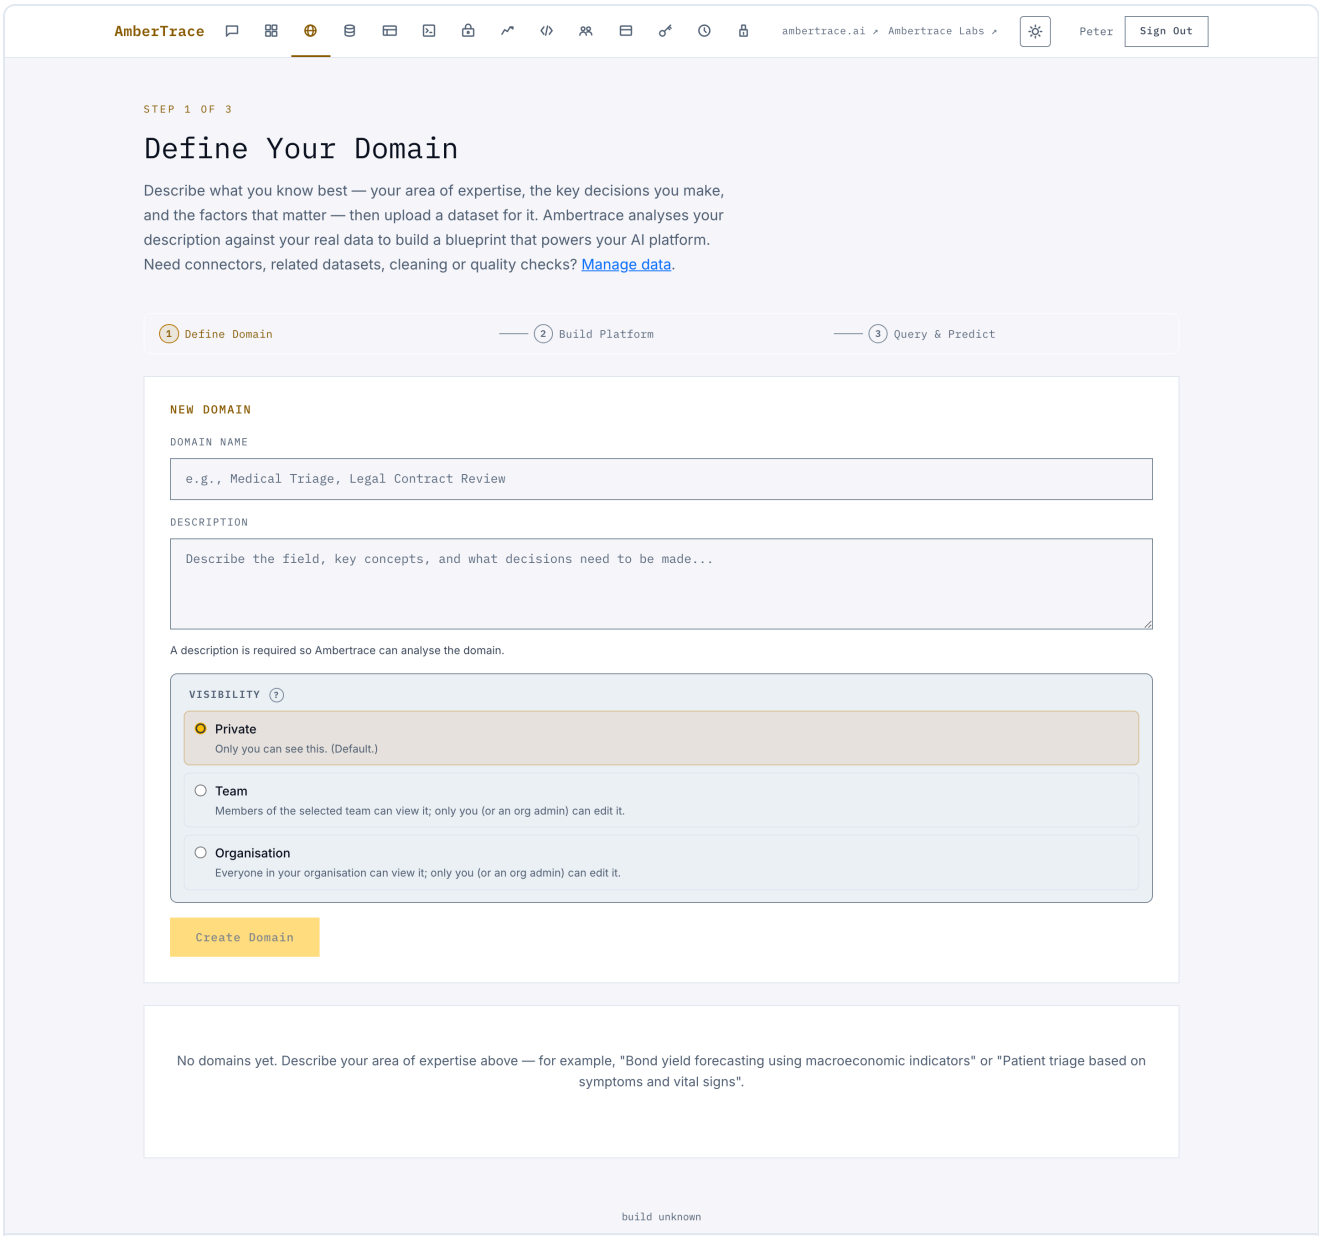

Domains — describe your field; AmberTrace structures it into a blueprint.

YOUR SCREENS, ONE BY ONE · STEP 2 OF 4

6. Data

Bring in your information. Upload a spreadsheet, or connect to a live data source, and link it to a domain.

What it's for

Your AI learns its patterns from your data. On this screen you upload a CSV or Excel file (or connect a live source) and tell AmberTrace which domain it belongs to. AmberTrace also checks the quality of your data and flags anything that needs attention.

When you'd use it

This is your **second step**, after defining a domain. Come back whenever you have fresh data to add.

Key actions

1. Choose the **domain** this data belongs to.
2. Pick a **data file** (CSV or Excel) to upload.
3. Optionally name the **Decision Column** — the column that records the outcome or decision (e.g. `outcome`, `label`) — if your data already has one.
4. Upload. AmberTrace produces a **quality report**: missing values by column and recommendations to clean things up.
5. Prefer a live feed? Use **Data Connectors** to pull in external series (for example, economic data from FRED, or market series) on a schedule.

Taxonomy Search

Use the **Taxonomy Search** panel to find the right connector or series before you configure one. Filter by asset class, country, region, or currency, or type a free-text query. Each result shows the connector type, series identifier, and a "Copy ID" button so you can paste the identifier straight into a connector configuration.

Data Connectors — the full catalogue

AmberTrace ships 24 built-in data connectors. Some require an API key you supply (noted below); keyless connectors pull public data directly. Where a source licence requires attribution, AmberTrace includes the notice automatically alongside every delivery of that data.

Connector	What it fetches	Key required?
FRED	Federal Reserve Economic Data — US macro time-series (GDP, CPI, rates, commodities, equity proxies).	Yes (free FRED key)
FRED Sentiment	Published sentiment indices from FRED (Michigan Consumer Sentiment, VIX, etc.).	Yes (free FRED key)
Yahoo Finance	Historical close prices for stocks and ETFs.	No

Connector	What it fetches	Key required?
Coinbase	Historical crypto close prices by trading pair.	No
SEC EDGAR	US company fundamentals (revenue, net income, assets) from XBRL filings.	No
US Treasury FiscalData	US debt levels, average interest rates, and exchange rates.	No
Bank of England	Gilt yields and SONIA historical series from the BoE Interactive Database.	No
BoE Yield Curves	Bank of England fitted yield curves (nominal, real, inflation, OIS) — full maturity grid.	No
Swap Curves (DTCC)	OIS swap-rate curves derived from DTCC public trade data for USD, EUR, GBP, JPY, CHF, and CAD.	No
ECB	Euro-area sovereign bond yield curves and other SDMX series from the ECB Data Portal.	No
IMF iData	Global macro data (CPI, exchange rates, balance of payments) via IMF SDMX dataflows.	Yes (IMF subscription key)
IMF WEO	World Economic Outlook annual CPI inflation forecasts (actual and projected) for curated countries.	No
Eurostat	EU macro, prices, labour, and trade statistics via SDMX.	No
OECD	OECD macroeconomic data (CLI, prices, activity) by dataflow.	No
ONS	UK CPI/CPIH and macro series by CDID code from the Office for National Statistics.	No
World Bank	Global development and macro indicators by country from World Bank Open Data.	No
EIA	US Energy Information Administration — electricity, petroleum, natural gas, coal, renewables, CO_2 emissions, and preset oil prices/production/inventories.	Yes (free EIA key)
CFTC	Commitments of Traders (COT) — weekly futures positioning data.	No
GDELT	Historical news sentiment (tone scores) from the GDELT global event database, back to 2015.	No
Sentiment (LLM)	Custom LLM-scored sentiment indices from RSS news feeds you specify, along dimensions you define.	No (but requires feed and dimension config)

Connector	What it fetches	Key required?
Polymarket	Historical traded probability for prediction-market outcomes.	No
Kalshi	Historical candlestick probability for prediction-market outcomes via the Kalshi Exchange API.	Yes (Kalshi key)
Betfair	Current runner implied probability from Betfair Exchange decimal odds.	Yes (Betfair credentials)
Generic REST/CSV	Fetch a table from any HTTP endpoint — bring your own URL and optional auth headers.	Optional (depends on endpoint)

Connector API keys are saved in **Settings** → **Account** (Section 13) and are stored securely per user.

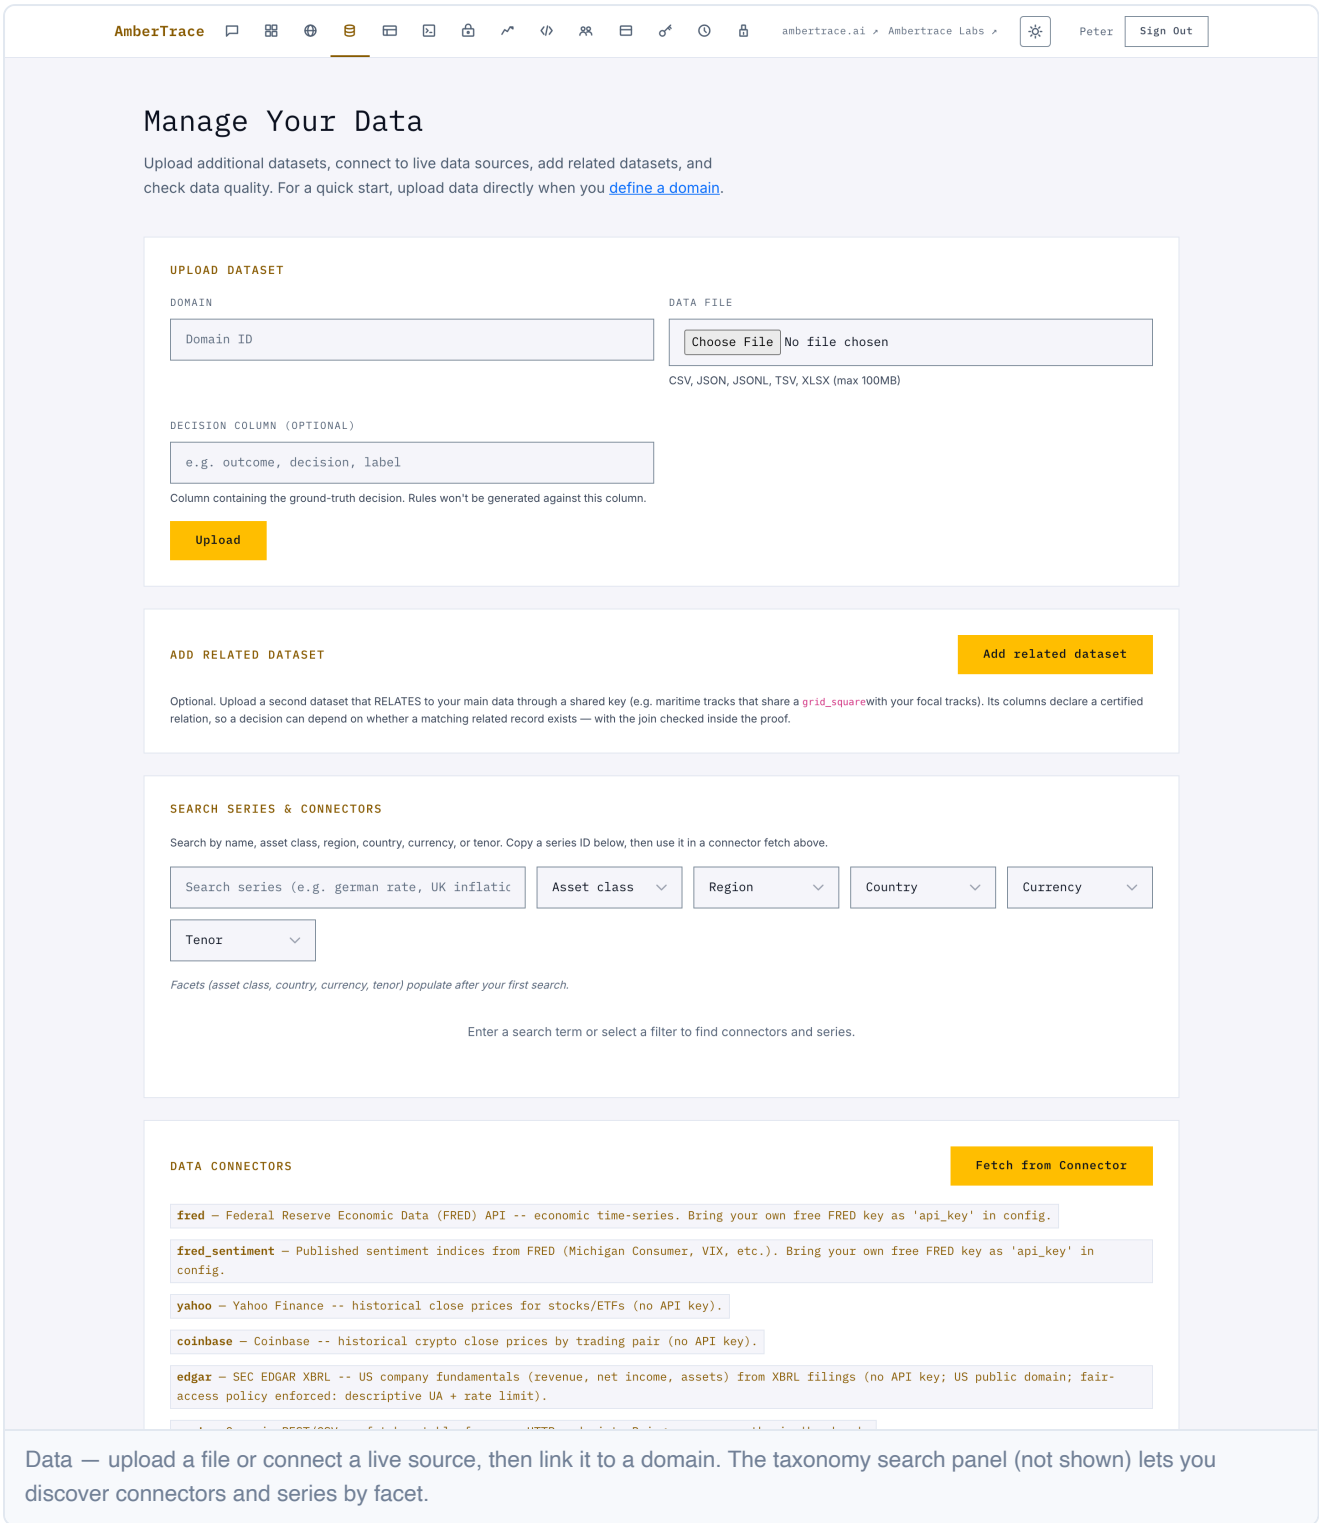

Data — upload a file or connect a live source, then link it to a domain. The taxonomy search panel (not shown) lets you discover connectors and series by facet.

YOUR SCREENS, ONE BY ONE · STEP 3 OF 4

7. Platforms

Build your AI. Combine your domain's rules with your data to create a working, explainable platform.

What it's for

A **platform** is your finished AI. Building one fuses the blueprint from your domain with the patterns in your data, so every answer it gives can be explained. You can watch the build happen and see a quality summary when it's done.

When you'd use it

This is your **third step**. Build a platform once your domain and its data are ready; rebuild when your rules or data change.

Key actions

1. Choose the **domain** to build from (it needs an active blueprint).
2. Optionally turn on **Verified profile** — this makes the platform's answers *provably* obey your rules (see the Console and Agent Policy Gate sections for why that matters).
3. Start the build and watch the **Build Progress** and log update live.
4. When finished, your platform appears in **Your Platforms**. Expand it to see its configuration and **sharing** options.

> **WHAT'S "VERIFIED PROFILE"?**
>
> A normal platform explains its answers. A *verified* platform goes further: it only returns an answer when it can prove the answer follows your rules. Choose it when a wrong answer would be costly — clinical, safety, or compliance decisions.

> **PLATFORM OR PREDICTION?**
>
> A **platform** answers questions about *what is true now*, from your data and rules — you ask it in the Console. To *forecast a value forward in time*, use **Predictions** instead. Section 16 makes the choice crisp.

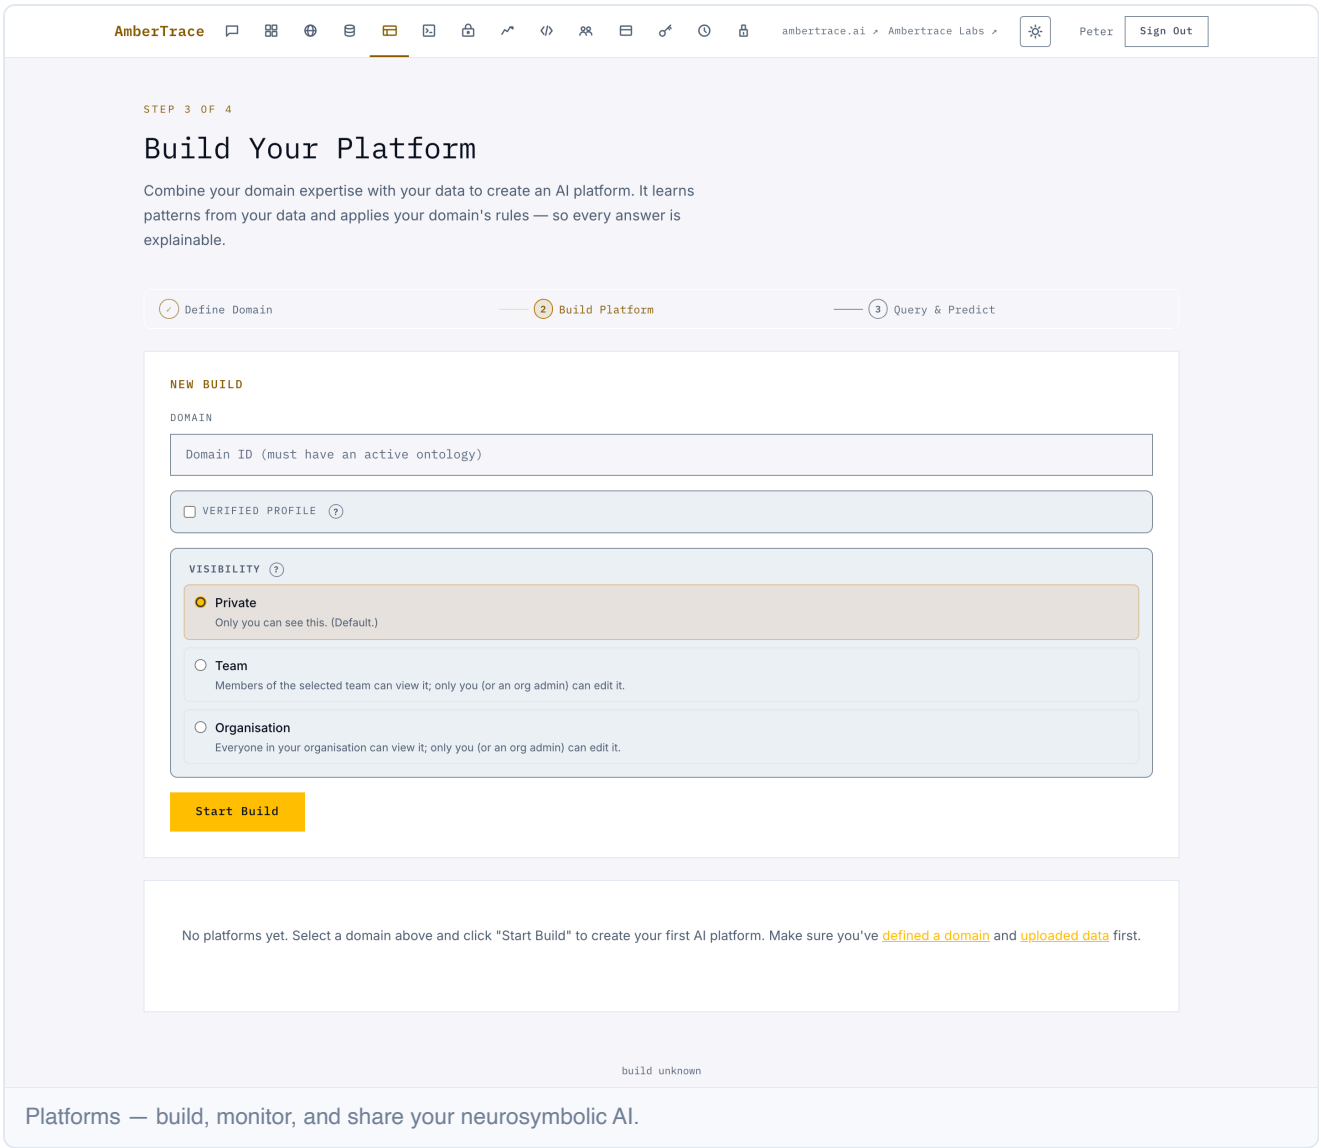

Platforms — build, monitor, and share your neurosymbolic AI.

YOUR SCREENS, ONE BY ONE · STEP 4 OF 4 · CAPABILITY-DEPENDENT

8. Console

Ask your platform a question and get an answer backed by your data and your rules — with the full reasoning shown.

> **CAPABILITY-DEPENDENT**
>
> The Console appears in the navigation only when your organisation's **query** capability is enabled. If you do not see it, contact your organisation administrator.

What it's for

The Console is where you **Query** your platform: type a question, get an answer, and see a complete explanation of how the AI reached it. For verified platforms you also get a **Verified Result** and an **Explanation Trace** showing the evidence and the rules that fired.

When you'd use it

This is your **fourth step** — and the one you'll use most. Any time you want a decision or an answer from your AI, ask it here.

Key actions

1. Select the **platform** you want to ask.
2. Type your **question** and submit.
3. Read the **Answer**, then open the **Explanation Trace** to see the neural evidence and the rules behind it.
4. Your **Query History** is kept — replay any past question with one click.

What a verified answer looks like — a sample Ambert Report

Every answer comes back as an **Ambert Report**: the answer itself, a confidence, the *neural evidence* that informed it, and the *rules that fired* — plus, for a verified platform, a machine-checked **proof**. Here is a worked example from a medical-triage platform.

> **AMBERT REPORT — MEDICAL TRIAGE PLATFORM**
>
> **Question:** "How should we handle a 71-year-old presenting with chest pain, troponin 0.9 ng/mL, and a prior cardiac event?"
>
> **Answer:** URGENT — escalate to cardiology · **Confidence:** 0.94 · **Verified:** ✓ proof attached

Part of the report	What it shows
Neural evidence	Elevated troponin and the patient's age band were flagged as high-risk signals, learned from patterns in your data (evidence weight 0.88).

Part of the report	What it shows
Rules that fired	`troponin ≥ 0.4 → cardiac_marker_positive` ; `cardiac_marker_positive ∧ prior_event → high_risk` ; `high_risk → URGENT`
Proof	Kernel-checked: the *URGENT* verdict provably follows from those rules given the evidence — no unproven step. On a verified platform, an answer that cannot be proven is not returned.

```
answer=URGENT · confidence=0.94 · verified=true · rules_fired=3 · proof=OK
```

REMEMBER THE DIFFERENCE

The Console answers questions about *what is true now*, given your data and rules. To *forecast the future*, use **Predictions**. Section 16 makes the distinction crisp.

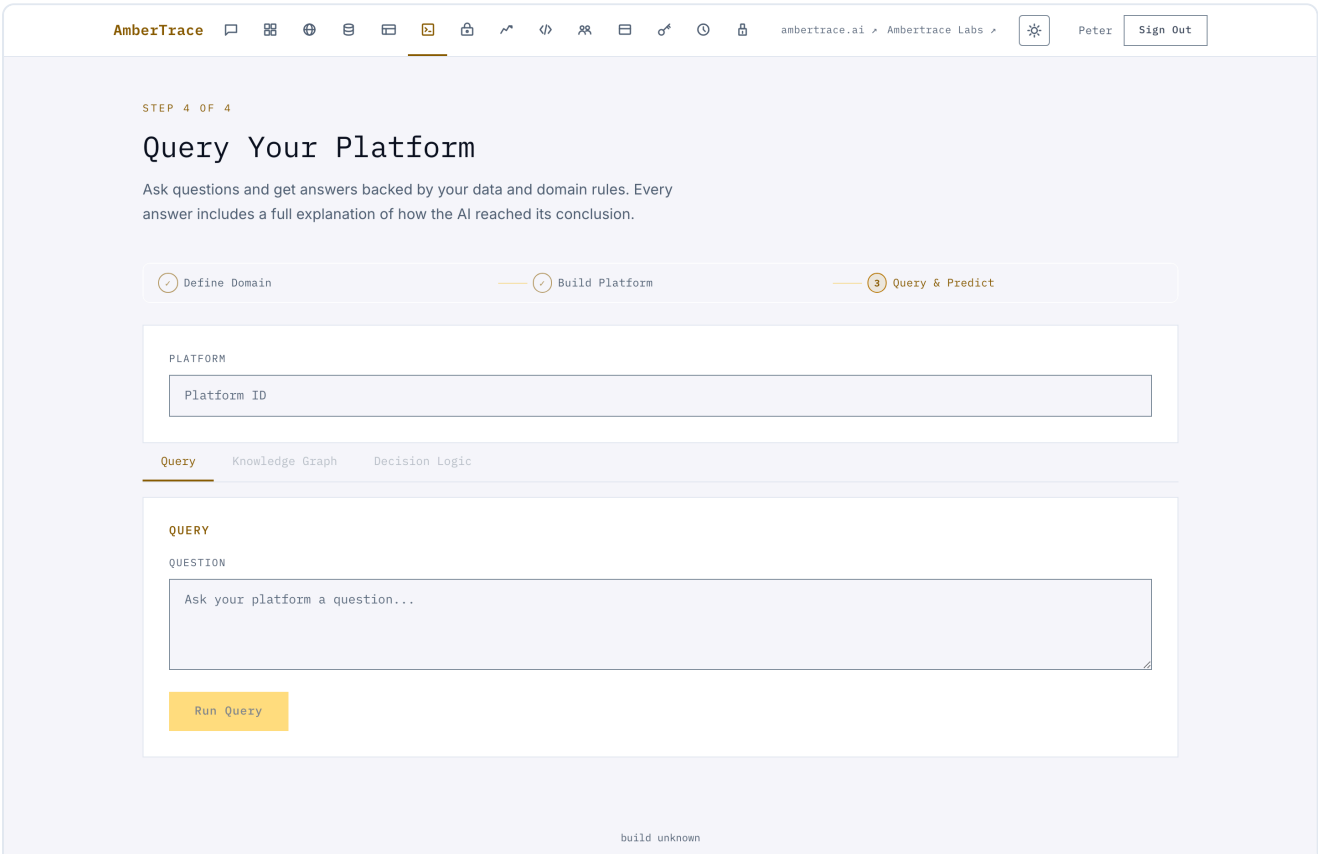

Console — ask a question; every answer comes with its explanation.

YOUR SCREENS, ONE BY ONE · SHOWN WHEN ENABLED

9. Agent Policy Gate

Let an AI agent act — but only within a policy you write in plain English. Nothing happens unless it can be *proven* allowed.

What it's for

This screen governs the *actions* an AI agent wants to take. You write a governance policy in plain English; AmberTrace turns it into verified rules; then every action the agent proposes is checked against that policy and gets a clear **PERMIT** or **DENY** — each permit carrying a machine-checked proof. It is **fail-closed**: if an action can't be proven to be within your policy, it does not happen.

> **THE TRUST BOUNDARY, IN ONE LINE**
>
> *The AI proposes; the checker proves.* AmberTrace is the sole executor — no action takes effect unless a kernel-checked proof shows it is inside your declared policy. Trust lives in the checker, never in the AI.

When you'd use it

Use it whenever an automated assistant or agent will *do* things on your behalf — approving, dispatching, releasing, escalating — and you need a guarantee it stays inside the rules. (This screen appears only when your organisation has the feature switched on.)

Key actions

1. Read the **Trust boundary** panel so you understand what's trusted and what isn't.
2. Write your policy under **Author policy (natural language)** — e.g. "Only approve a payment if a second person has signed it off."
3. Generate it. AmberTrace shows which rules were **admitted** and which were **rejected** (and why), so nothing slips in silently.
4. Watch proposed actions get a **PERMIT** (with proof) or **DENY** verdict against your policy.

ARIA · VERIFIED PROFILE

Agent Policy Gate

Plug in an (untrusted) AI agent and watch every action it proposes pass — or fail — a verified policy you author in plain English. Ambertrace is the sole executor: **no action takes effect unless a kernel-checked proof shows it is within your declared policy.** Fail-closed. The LLM proposes; the kernel proves — structural containment is proven, not agent wisdom.

TRUST BOUNDARY

UNTRUSTED
Agent (LLM)
Proposes actions (tool calls). Its reasoning is never relied upon — quality is irrelevant to the guarantee.

→

VERIFIED
Ambertrace gate
Derives the permit decision over your declared controls and checks the proof in the trusted kernel. Sole executor.

→

EFFECT
Action effect
A mock effect. Reached only after a certified permit with a machine-checked proof.

EXAMPLE POLICIES

One verified core, many domains. Select an example to load its plain-English policy into the author panel and its sample actions into the test panel.

COMPLEX HEALTHCARE DELIVERY & TRIAGE

Clinical triage assistant
Escalate to a clinician any message that mentions self-harm. Do not approve an opioid refill unless prior_authorization is present. Only allow actions of type triage, schedule, or refer.

ENERGY-SYSTEM DISPATCH / GRID

Grid dispatch controller
Reject any dispatch that pushes line_loading_pct over 90. Only dispatch generators in the approved set (GEN_A, GEN_B, GEN_C). Block any dispatch over 500 MW unless stability_check_passed is true.

EMBEDDED & CYBER-PHYSICAL CONTROL (OPERATIONAL-AUTOMATION SAFETY)

Actuator command guard
Block any actuator command with pressure outside 2 to 8 bar. No command over 100 units. Reject commands flagged is_runaway.

RISK ASSESSMENT & MANAGEMENT (ACCESS PDP)

Privileged access decision point
Deny access to the OT network from an unmanaged device. Require mfa_passed for privileged requests. Only permit when change_ticket_approved is present.

SOFTWARE SUPPLY-CHAIN ASSURANCE

Deployment supply-chain gate
Block any deployment using an unsigned dependency. No action outside the approved namespace (prod, staging). Require code_review_passed for production deploys.

SOFTWARE SUPPLY-CHAIN ASSURANCE (CI/CD MERGE CONTROL)

Review-before-merge gate
Permit a merge into a protected branch only when required_reviews_passed is present, the approvals_count is at least 2, and ci_status is green. Block the merge otherwise.

SOFTWARE SUPPLY-CHAIN ASSURANCE (CI/CD RELEASE CONTROL)

Change-window release gate
Permit a production release only when within_change_window is true and change_ticket_approved is present. Block any release outside the approved change window or without an approved change ticket.

Agent Policy Gate — plain-English policy in; proven permit/deny out.

YOUR SCREENS, ONE BY ONE · CAPABILITY-DEPENDENT

10. Predictions

Forecast values in your data — and understand *what drives* each forecast, in plain rules you can read.

> **CAPABILITY-DEPENDENT**
>
> Predictions appears in the navigation only when your organisation's **predictions** capability is enabled. If you do not see it, contact your organisation administrator.

What it's for

Predictions looks *forward*. You pick something to forecast (a number that changes over time), train a model, and see how accurate it is — with explanations of the factors driving each forecast. AmberTrace's headline approach is **Symbolic**: it fits transparent driver-rules to your data, so you can read the rules and see the fit, not just trust a black box.

When you'd use it

Use Predictions for questions about the *future*: next month's demand, a series' likely trend, an economic indicator's path. (For questions about the present, use the Console — see Section 16.)

Key actions

1. Choose your **forecast approach**. **Symbolic (explainable rules)** is the recommended default; **Statistical** and **Neurosymbolic** options are available too.
2. Configure what to predict: pick the platform/data and the field to forecast.
3. Choose an **Optimisation Objective** (time-series mode only — see below).
4. If using the Statistical or Neurosymbolic approach, choose a **Model Type** (see below).
5. **Train** the model, then review its accuracy (how well it fits, and how much skill it adds over simply assuming "tomorrow looks like today").
6. Open the **Predict** view to see the forecast, the chart, the discovered driver-rules, and the explanation.

Optimisation Objective

For time-series predictions you can choose what the model optimises for. The default — **Skill vs persistence** — measures how much the model improves on the naïve assumption that "tomorrow looks like today." Three trading-oriented alternatives are also available:

Objective	What it optimises
Skill vs persistence (default)	Forecast accuracy relative to the persistence baseline — how much better is the model than "no change"?
Directional PnL	Cumulative profit-and-loss from following the model's predicted direction of change.

Objective	What it optimises
Sharpe ratio	Risk-adjusted return: annualised mean PnL divided by its standard deviation.
Hit rate	Fraction of periods where the model correctly predicts the direction of the next move.

> **FREQUENCY IS REQUIRED FOR TRADING OBJECTIVES**
> The three trading objectives (*Directional PnL*, *Sharpe ratio*, *Hit rate*) need a declared data frequency (daily, weekly, monthly, or quarterly) in order to annualise returns correctly. If your prediction config uses cross-sectional mode (no time axis), only *Skill vs persistence* is available.

Trading metrics in the forecast response

When you train with any objective, AmberTrace's symbolic-forecast response includes per-objective trading metrics alongside the forecast value and explanation. You will see:

- **directional_pnl** — cumulative PnL from following the predicted direction.
- **hit_rate** — fraction of periods with a correct directional call.
- **sharpe_ratio** — annualised risk-adjusted return (shown when enough data points are available).
- **max_drawdown** — largest peak-to-trough decline in cumulative PnL (a risk measure, always reported, never used as the optimisation target).

These metrics appear in the API response metadata and on the Predictions screen's results panel, giving you a clear picture of both forecast quality and trading viability.

Neural model selection (Statistical & Neurosymbolic approaches)

When you choose the **Statistical** or **Neurosymbolic** approach, a **Model Type** dropdown lets you pick the underlying neural learner:

Model	What it is
GBT (default)	Gradient Boosted Trees — fast, robust, and the best starting point for most tabular data.
Ridge	Ridge Regression — a simple linear model with regularisation; good when relationships are roughly linear.
Lasso	Lasso Regression — like Ridge but drives unimportant feature weights to zero, acting as a built-in feature selector.
LSTM	Long Short-Term Memory network — a recurrent neural network suited to longer sequential patterns. Slower to train.
Transformer	Transformer architecture — attention-based, strong on complex temporal dependencies. Slower to train.

In **Symbolic** mode (the recommended default), the model type is handled automatically — the symbolic forecaster fits its own transparent driver-rules and does not expose a model-type choice.

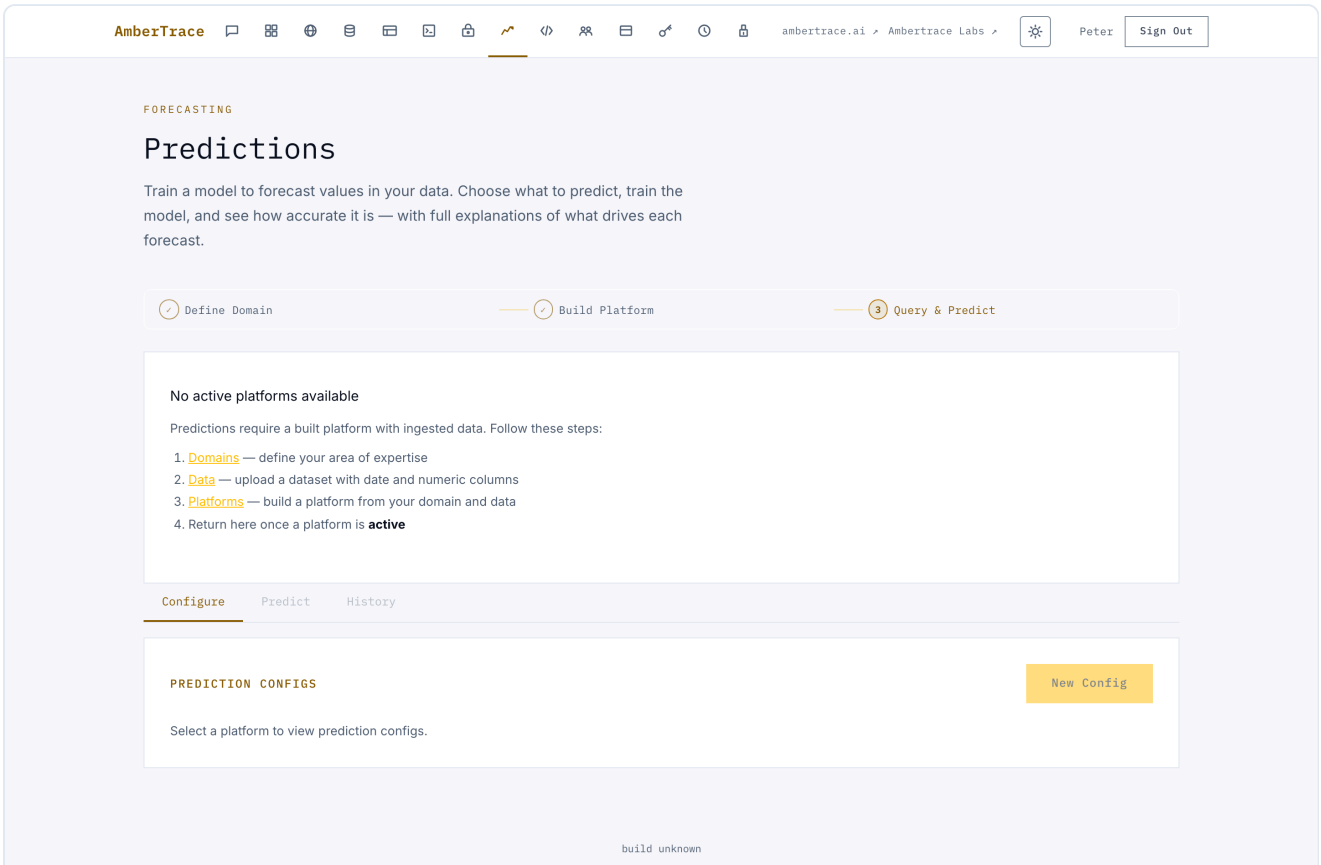

Predictions — forecast values and read the rules that drive them. The objective picker and model-type selector appear during configuration.

YOUR SCREENS, ONE BY ONE

11. API

Create keys so software — including your own AI assistant — can use your platforms safely. (Full walkthrough in Section 17.)

What it’s for

The API screen issues **API keys**: secret passcodes that let a program talk to your platforms without logging in by hand. There are two kinds — **Agent keys** (full access, like your own account) and **Platform keys** (query-only, tied to one platform).

When you’d use it

Use it when you want to connect AmberTrace to another tool, or hand a key to an AI assistant so it can ask your platform questions on your behalf.

Key actions

1. Choose **Agent Key** or **Platform Key**.
2. Name the key (e.g. “Claude Agent”), pick the platform if it’s a platform key, and click **Generate Key**.
3. **Copy the key immediately** — it’s shown only once.
4. Manage existing keys in the table, and **Revoke** any you no longer trust.

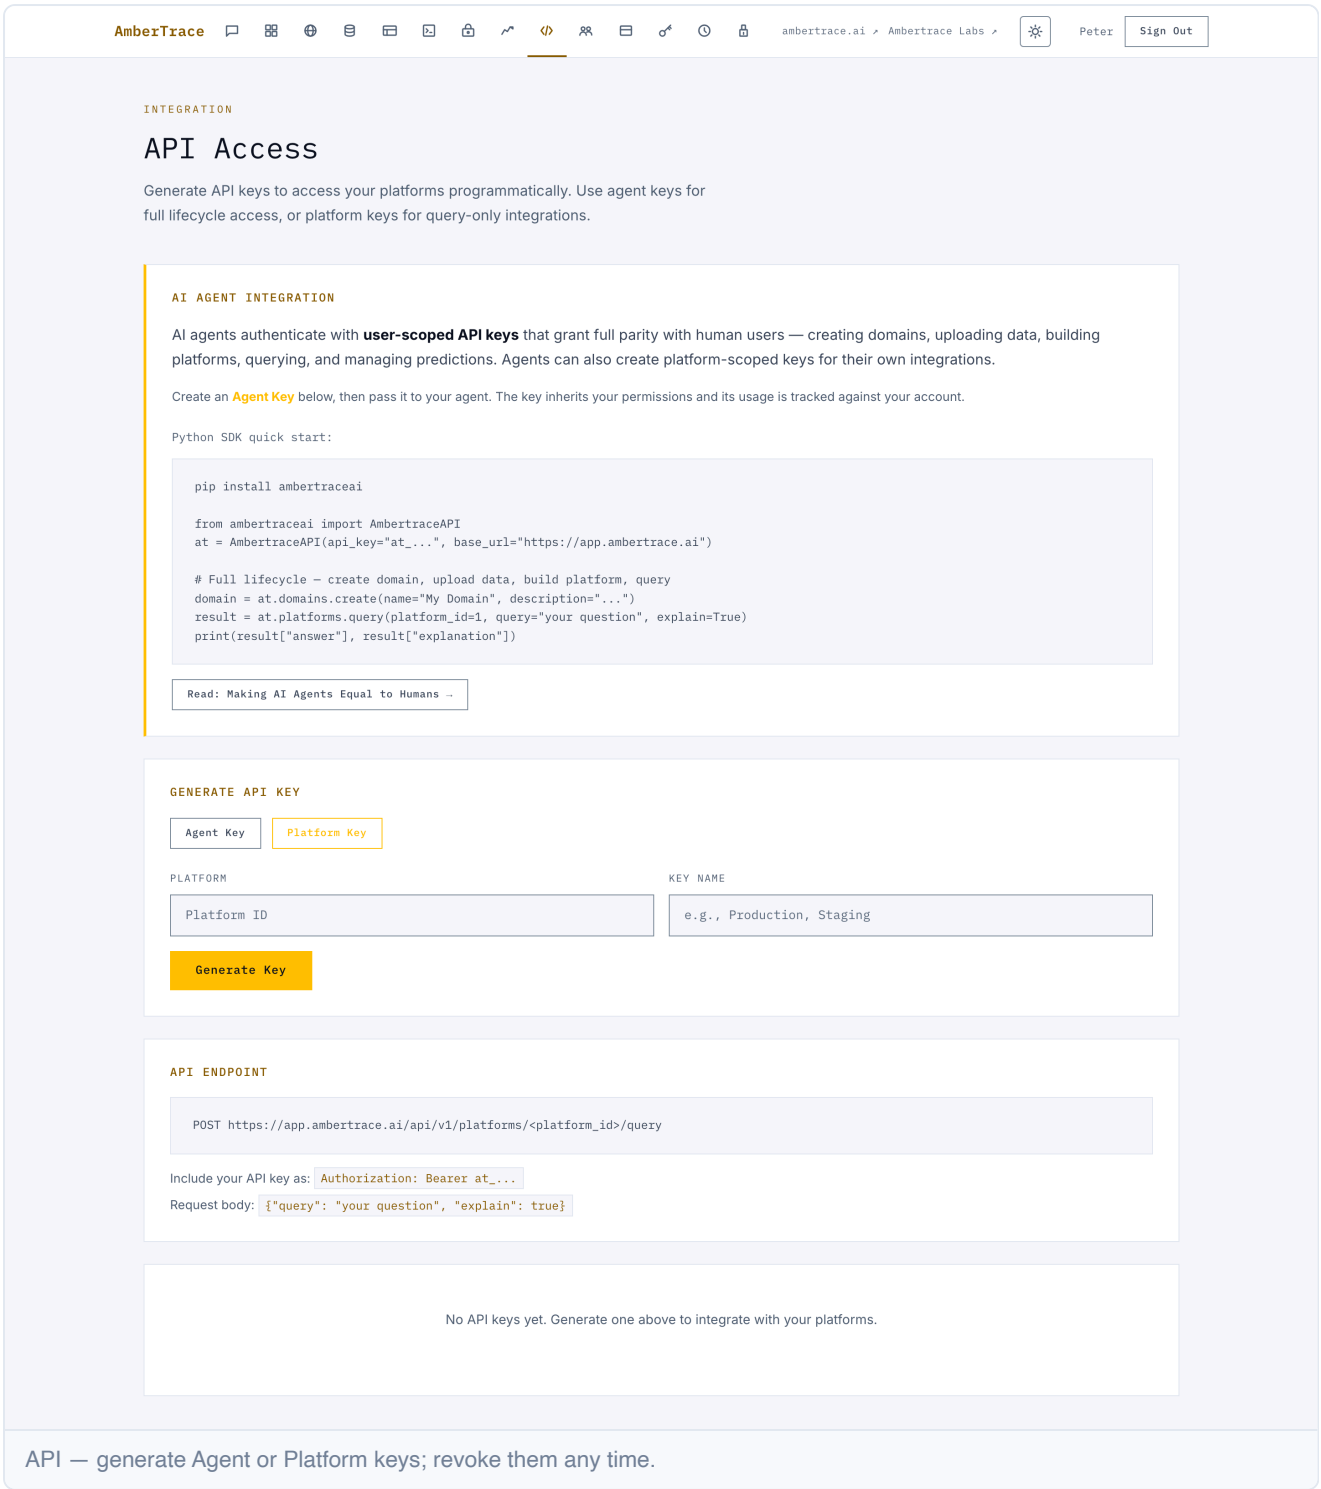

API — generate Agent or Platform keys; revoke them any time.

YOUR SCREENS, ONE BY ONE · ORGANISATION ADMINS

12. Teams

Share your work with colleagues — safely, and only inside your own organisation.

What it's for

Teams let you group colleagues and share a domain or platform with all of them at once. Shared members get **view access**; editing rights stay with the owner. There are three sharing levels for any domain or platform:

- **Private** — visible only to you (the default).
- **Team** — visible to the members of a team you belong to.
- **Organisation** — visible to everyone in your organisation.

Who you can share with

Sharing is **strictly within your own organisation**. Members are added by an **organisation administrator**, who selects from people who *already have an account in your organisation* — you cannot add an arbitrary external email address, and no one outside your organisation can ever see your work. By design there are **no public share-links**: we chose not to allow link-based sharing, for security. (Link-based or external sharing can be enabled for enterprise customers on request.)

How people are notified

When you are added to a team, when a domain or platform is shared with you, and when access is later **revoked**, you receive a notification (in-app and by email, according to your notification settings). Shared items also appear in your **Domains** / **Platforms** lists.

Key actions

1. An admin creates a **New Team** and adds members from your organisation.
2. From a domain or platform, set its **sharing level** to Team (choose a team) or Organisation.
3. Change the level back to **Private** at any time to revoke access — affected members are notified.

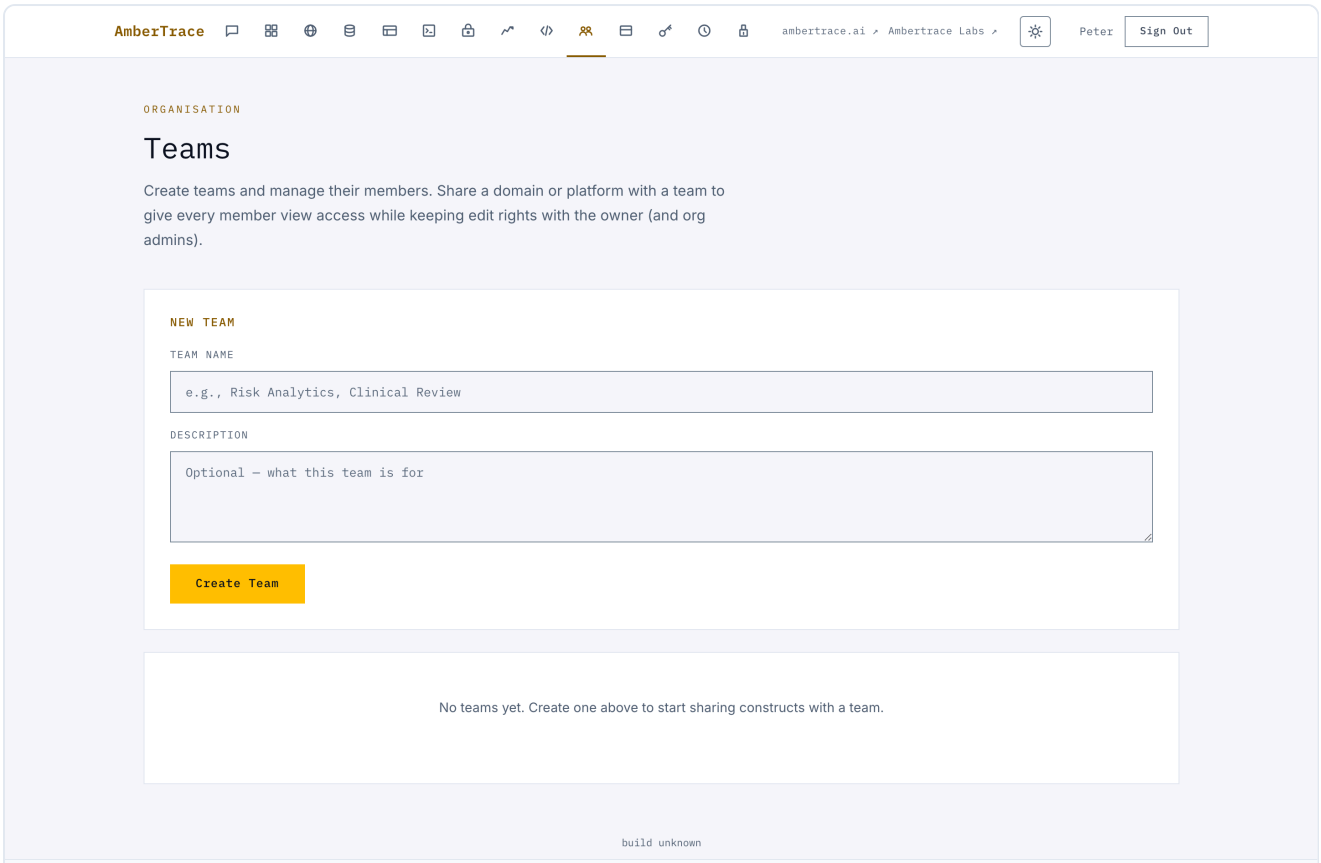

Teams — group colleagues and share view access, only within your organisation.

YOUR SCREENS, ONE BY ONE

13. Settings

Your account details, security settings, and — for organisation administrators — the audit log and SSO configuration.

Settings is reached by clicking your name in the top-right corner. It is organised into tabs:

Account

Review your username, email, and organisation. Save or update **connector API keys** for live data sources (such as a FRED key for economic data).

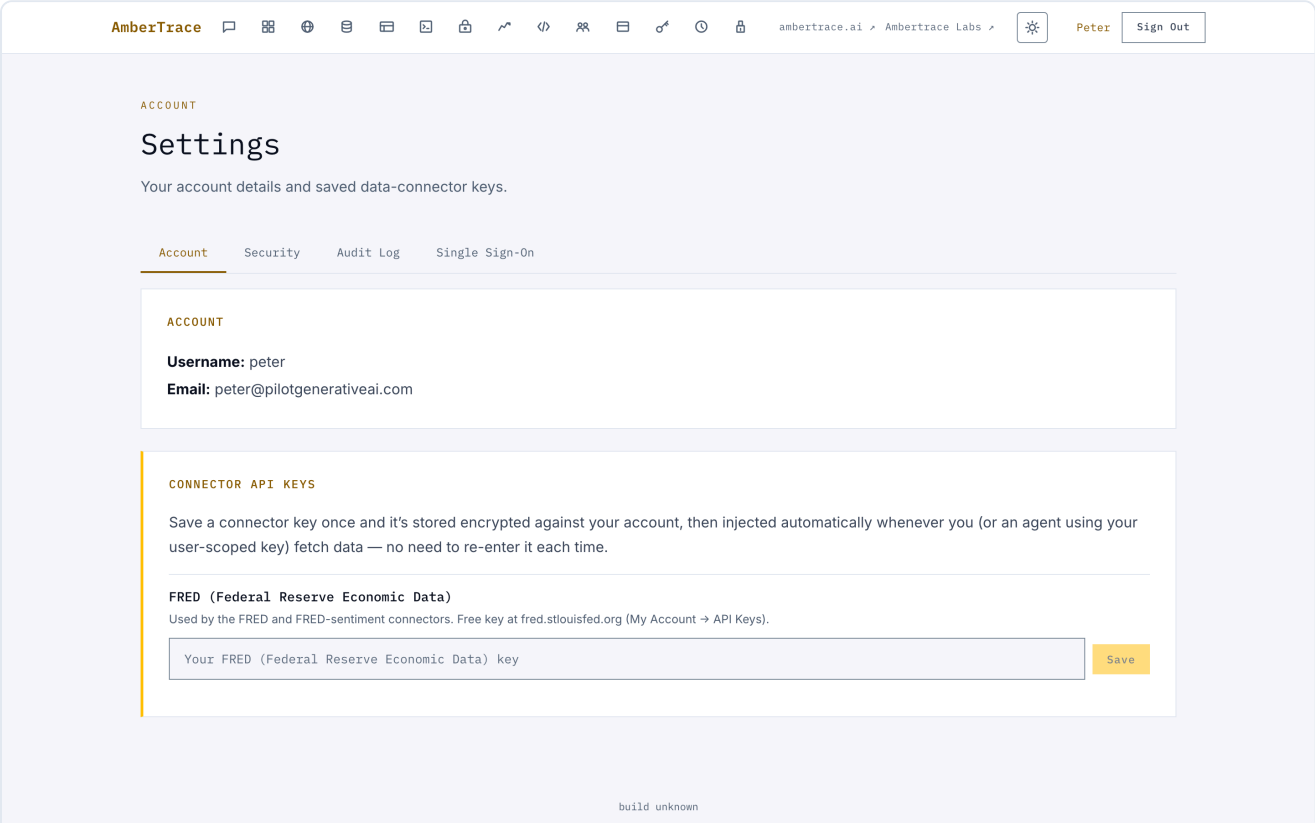

Settings — Account tab: your details and saved connector keys.

Security (MFA)

Set up **multi-factor authentication** (MFA) for your account using a TOTP authenticator app (Google Authenticator, Authy, 1Password, or similar). Once enrolled, you enter a one-time code at each login in addition to your password.

Organisation administrators also see an **Organisation MFA Policy** panel here. Enabling it requires *all* password-login members of the organisation to use MFA — members who have not yet enrolled will be prompted to set up MFA at their next sign-in. Members who sign in via SSO are unaffected (their identity provider manages MFA).

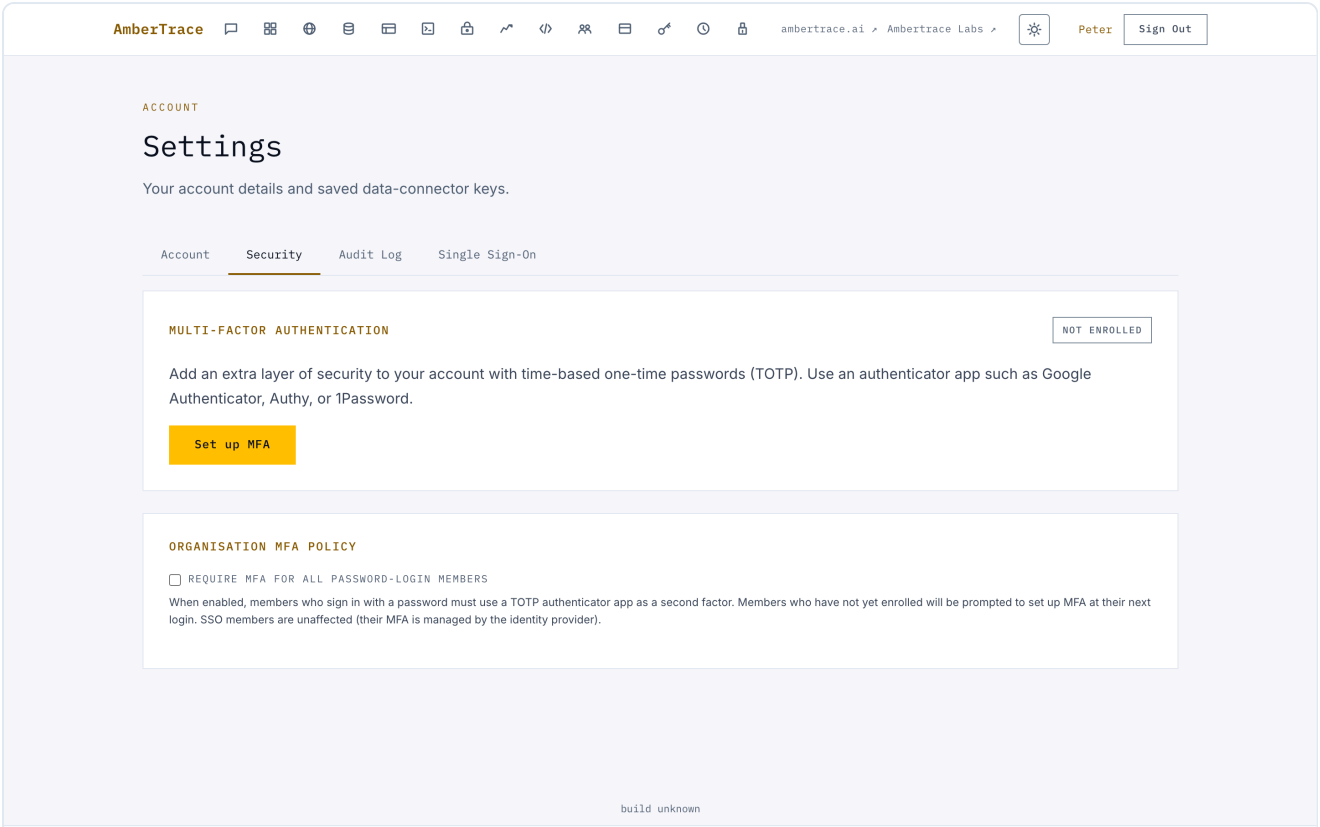

Settings — Security tab: MFA enrollment and organisation-wide MFA policy.

Audit Log Organisation admins

A complete, append-only record of sharing and team-membership changes across your organisation: who was given access, what was shared, what was revoked, and when. Use it for compliance reviews or to investigate access changes.

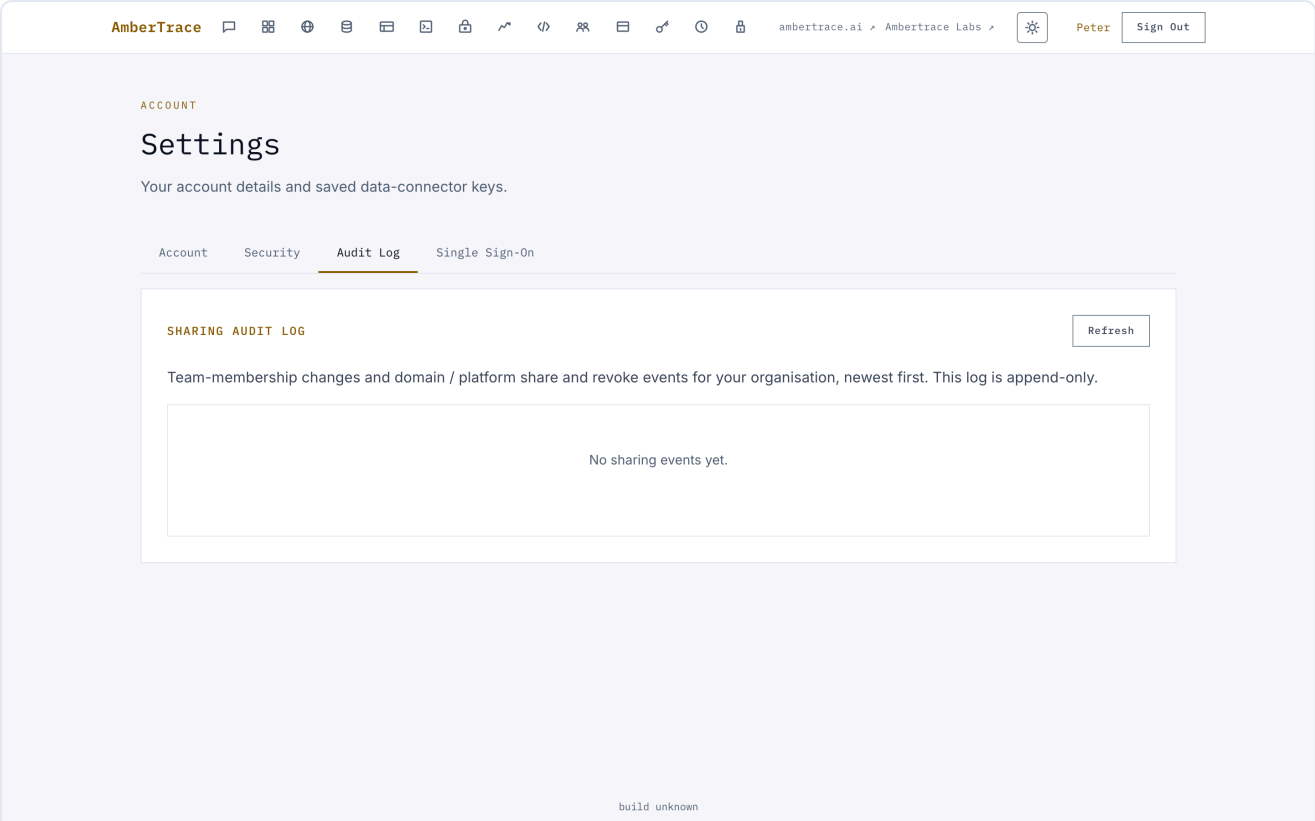

Settings — Audit Log: every share, revoke, and team change, newest first.

Single Sign-On (SSO) Organisation admins

Configure your organisation’s SSO identity provider so members can sign in with their corporate credentials instead of an AmberTrace password. AmberTrace supports OpenID Connect (OIDC) — enter your provider’s issuer URL, client ID, and client secret, and AmberTrace handles the rest.

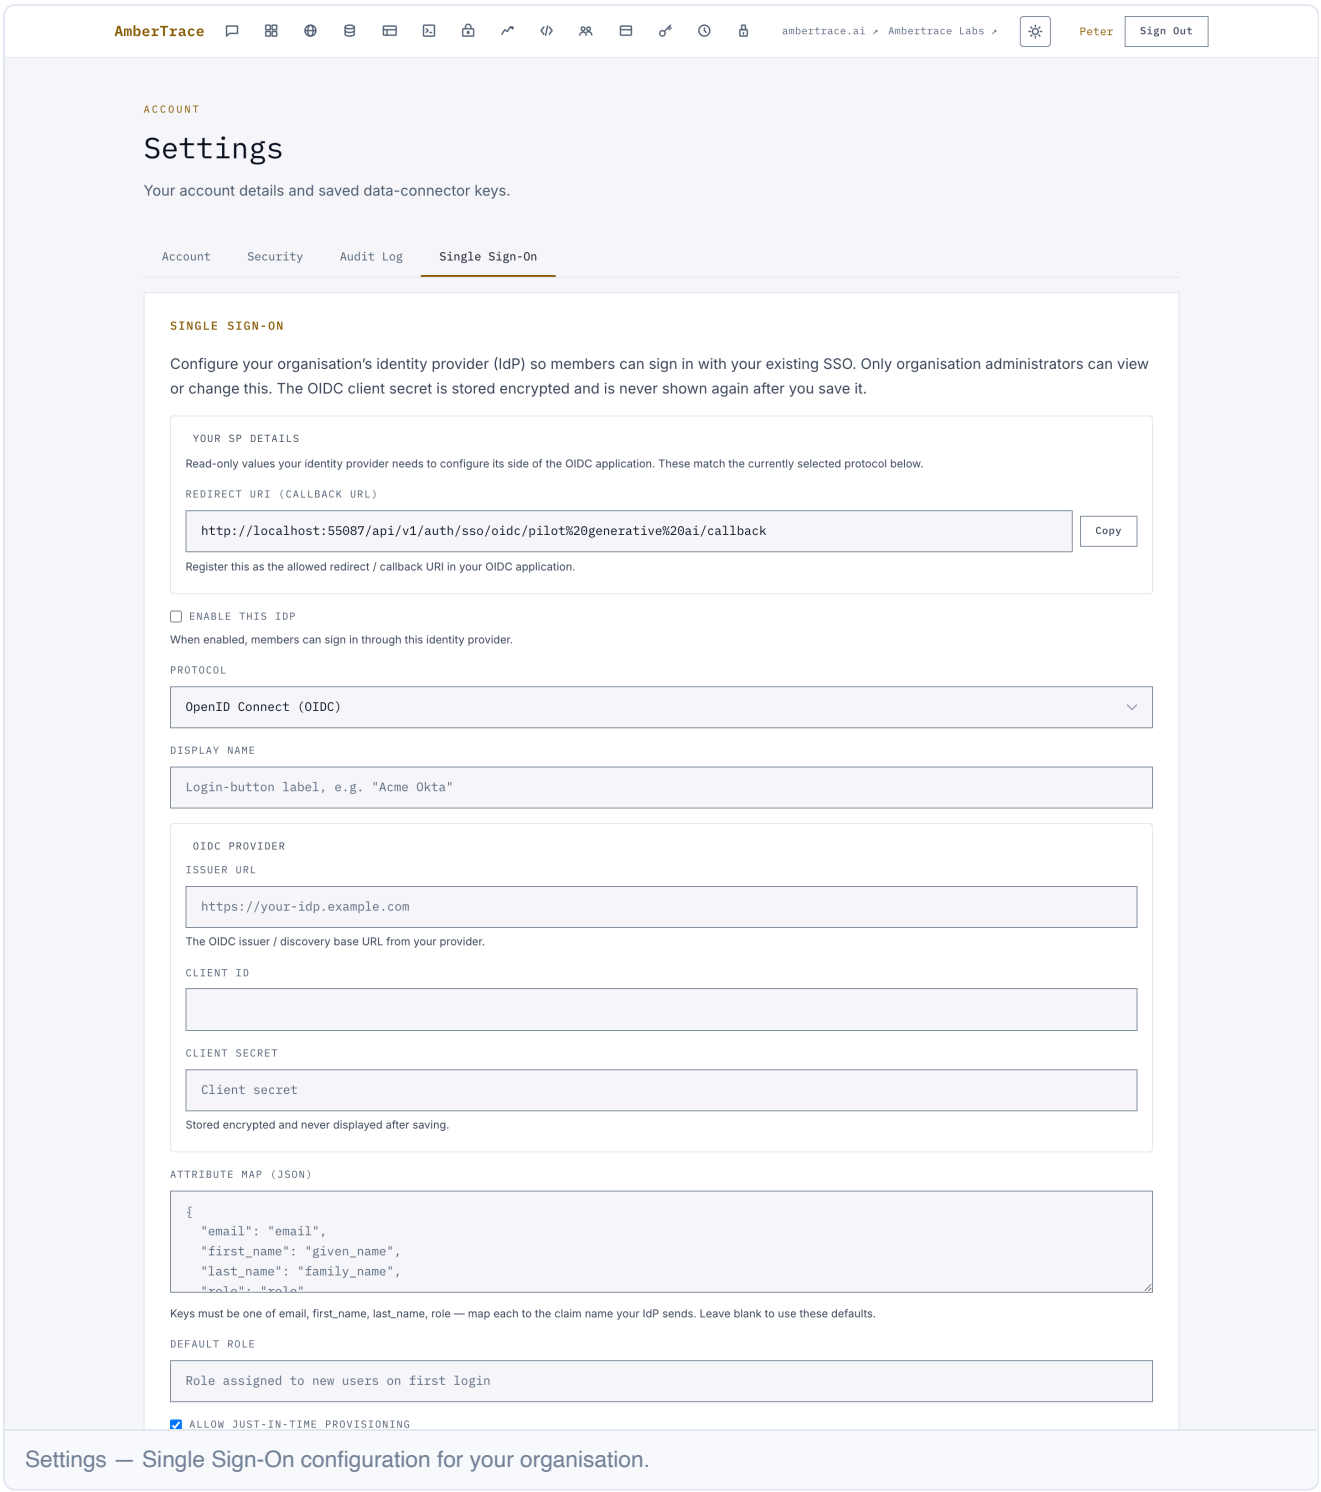

Settings — Single Sign-On configuration for your organisation.

ORGANISATION ADMINISTRATION · ORGANISATION ADMINS

14. Billing

Manage your organisation's subscription, seat count, and payment details.

What it's for

The Billing screen shows your organisation's current subscription status, how many seats are in use, and your entitlement. Payment is handled securely on Stripe's hosted pages — card details never touch AmberTrace.

When you'd use it

Check in here to review your seat usage, manage your payment method, or confirm your subscription tier. This screen is visible only to **organisation administrators**.

Key actions

- Review your **current status**: subscription tier and seats in use vs entitlement.
- Click **Refresh** to sync with the latest billing state.
- Manage your **payment method** via the secure Stripe portal (opens in a new tab).

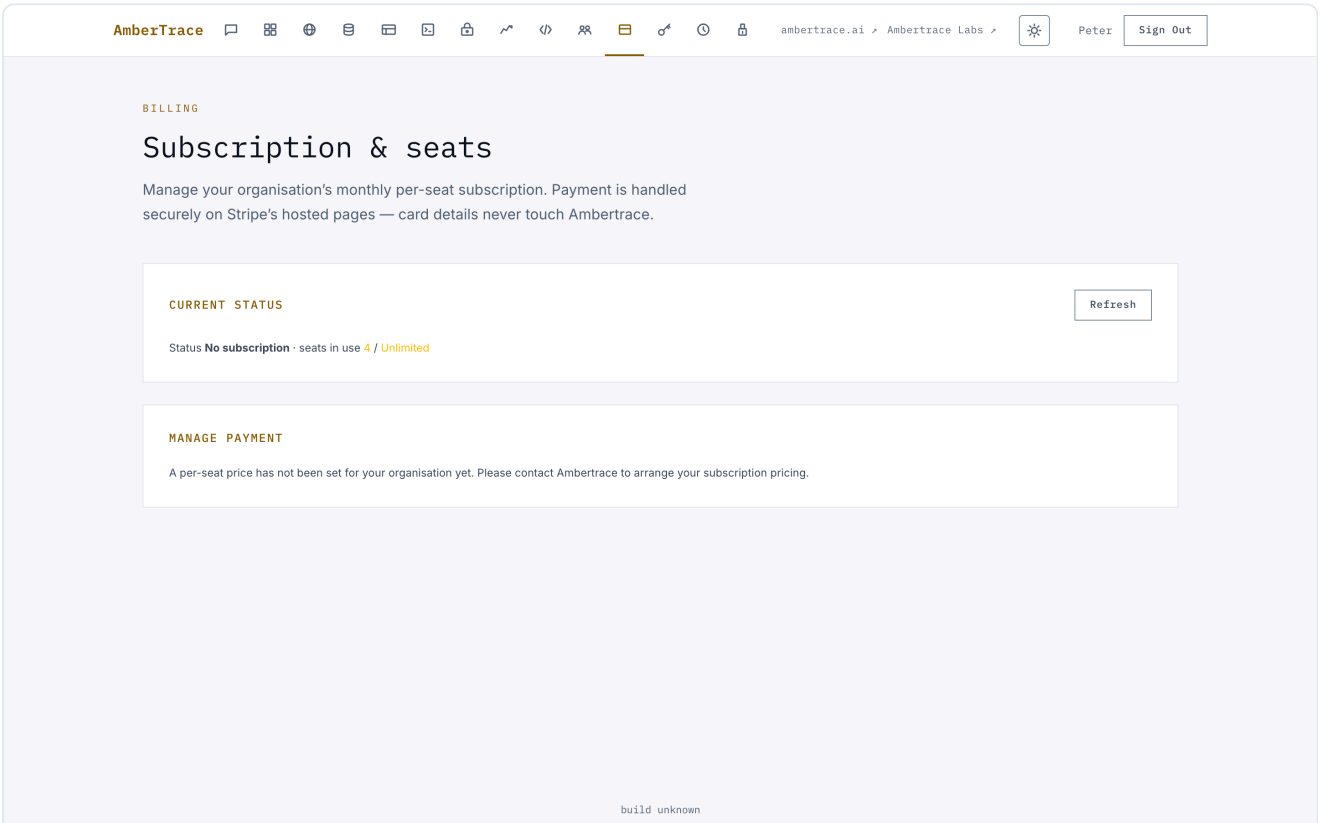

Billing — subscription status, seat usage, and payment management.

ORGANISATION ADMINISTRATION · ORGANISATION ADMINS

15. Encryption Keys

Manage your organisation’s per-tenant data-encryption key — status, rotation, and the revoke kill-switch.

What it’s for

AmberTrace encrypts your organisation’s data with a dedicated encryption key. This screen lets organisation administrators view the key’s status, rotate it on a schedule, and — in a security incident — revoke it immediately to render all encrypted data unreadable until it is restored.

When you’d use it

Use it for routine key rotation (a security best practice) or, in an emergency, to instantly cut off all reads and writes to your organisation’s encrypted data. This screen is visible only to **organisation administrators**.

Key actions

- Review **Key Status**: the active version, creation date, and last rotation.
- **Rotate key** to generate a new version — the old version is retired gracefully.
- In the **Danger zone**, **Revoke key** immediately stops all decryption for your organisation. Use only in a confirmed security incident.
- The **Key Audit Log** at the bottom records every lifecycle event (created, rotated, revoked, restored), append-only, with no key material ever shown.

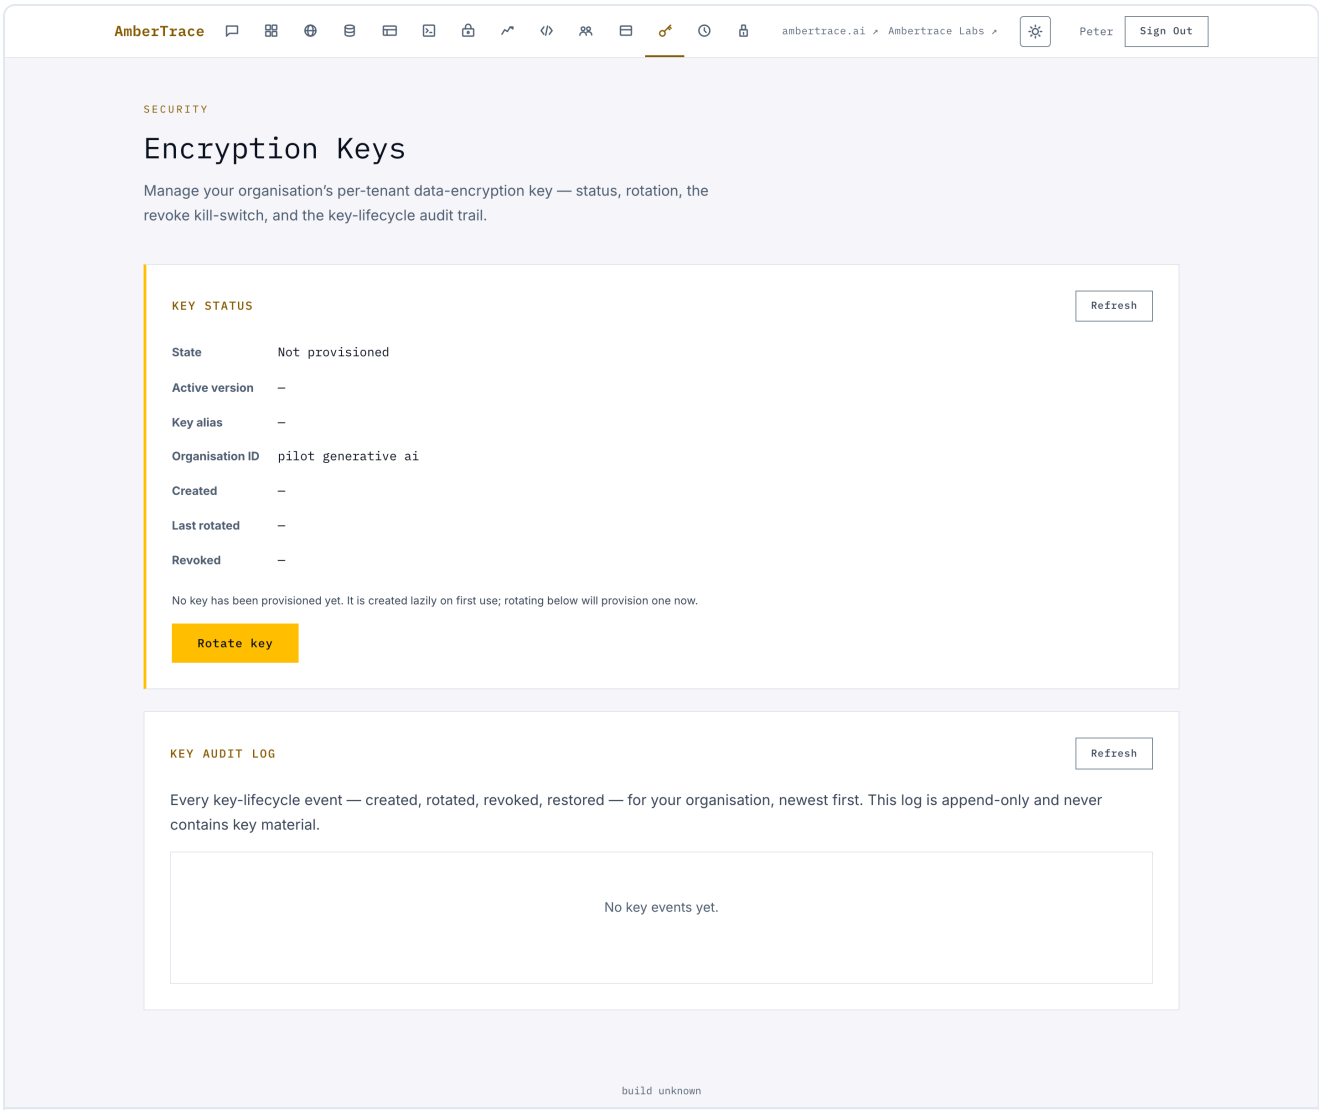

Encryption Keys — key status, rotation, revoke kill-switch, and audit trail.

THE TWO BIG TOOLS, EXPLAINED

16. Query vs Predictions — which one do I need?

AmberTrace gives you two ways to get answers. They sound similar but do very different jobs. This is the one distinction worth getting right.

Query Console

"Given what I know, what's the right answer now?"

You ask your built platform a question. It answers using your data *and* your rules, and shows you exactly how it got there — an explainable, **verified** decision or answer.

- Looks at the **present**: classify, decide, assess, permit/deny.
- Backed by an **explanation** — and, for verified platforms, a **proof** that it obeys your rules.
- Includes the **Agent Policy Gate**: permit/deny an AI agent's action, with proof.

Predictions Predictions

"Where is this number heading?"

You train a model on values that change over time and project them **forward**. AmberTrace fits transparent driver-rules to your data, so you can read what drives the forecast — not just trust a number.

- Looks at the **future**: forecast a trend, a demand level, an indicator.
- Time-series and macro-economic **forecasting**.
- Choose your **objective**: accuracy vs persistence, directional PnL, Sharpe ratio, or hit rate (Section 10).
- Forecasts can themselves feed a verified decision (forecast → fact → rule-checked signal).

A simple rule of thumb

If your question is…	Use	Where
"Should this referral be treated as urgent?"	**Query**	Console
"Does this case meet our approval criteria?"	**Query**	Console
"Is this agent's proposed action allowed by policy?"	**Query**	Agent Policy Gate
"What will demand look like next quarter?"	**Predictions**	Predictions
"Which way is this indicator trending?"	**Predictions**	Predictions

THE SHORT VERSION

Query = ask your platform a question and get an explainable, verified answer or decision about your domain *right now*. **Predictions** = forecast values that change over time, with readable rules

> explaining the forecast. Decide *now* → Query. Forecast *forward* → Predictions.

THE TWO BIG TOOLS, EXPLAINED

17. Giving your AI assistant a key (safely)

You can let an AI assistant use your AmberTrace platform on your behalf — asking it questions, even building things — by giving it an **API key**. Here's how, and why it's safe.

The idea, in plain terms

An **API key** is like a named pass card for software. Instead of your assistant logging in with your password, you hand it a key. Software then uses that key to talk to your platforms directly. You can see when each key was last used, and you can **revoke** (cancel) any key instantly — without changing your own password.

Two kinds of key — choose the right one

Key type	What it can do	Best for
Agent Key full access	Everything your own account can do: create domains, upload data, build platforms, manage rules, query, manage predictions. (It cannot create *other* agent keys.)	An AI assistant acting as a full member of your team.
Platform Key query-only	Talks to *one* specific platform — typically just to ask it questions.	A narrow integration that only needs to query a single platform.

> PICK THE LEAST POWERFUL KEY THAT DOES THE JOB
>
> If your assistant only needs to *ask questions* of one platform, give it a **Platform Key**. Reserve **Agent Keys** for assistants you trust to act with the full range of your own permissions.

Step by step: create a key for your assistant

1. Go to the **API** screen (the icon that looks like angle brackets).
2. Choose **Agent Key** (full access) or **Platform Key** (one platform, query-only). For a platform key, pick the platform.
3. Give the key a clear **name** so you'll recognise it later — e.g. "Claude Agent" or "Reporting Bot."
4. Click **Generate Key**. A long secret starting with `at_` appears.
5. **Copy it now** and paste it into your assistant's configuration. AmberTrace shows the full key only once; afterwards you'll only ever see the first few characters.

> KEEP IT SECRET
>
> Treat an API key like a password. Don't email it or paste it into shared documents. If a key is ever exposed, open the API screen and click **Revoke** — it stops working immediately, and you can issue a fresh one.

How your assistant uses the key

Once your assistant has the key, it asks your platform questions over the internet. In practice that looks like the small example shown on the API screen:

```
pip install ambertraceai

from ambertraceai import AmbertraceAPI
at = AmbertraceAPI(api_key="at_...", base_url="https://app.ambertrace.ai")

# Ask a platform a question, and get the explanation back too
result = at.platforms.query(platform_id=1, query="your question", explain=True)
print(result["answer"], result["explanation"])
```

You don’t have to write this yourself — your assistant’s developer or the assistant itself handles it. What matters for you: the key carries *your* permissions, and its use is tracked against your account.

Why it’s safe: the agent is governed

Handing an AI a key doesn’t mean handing over control. When your platform is a **verified** one — and especially behind the **Agent Policy Gate** (Section 9) — every action the assistant proposes is checked against the rules *you* wrote. The assistant *proposes*; AmberTrace *proves*. If an action can’t be proven to be within your policy, it is refused. So you get the convenience of an AI assistant *and* the guarantee that it stays inside your guard-rails.

> **THE WHOLE STORY IN ONE SENTENCE**
>
> Give your assistant a key so it can ask your platform questions and act for you — and rest easy, because a verified platform (and the Agent Policy Gate) will only let through what your own rules provably allow.

TRUST & SECURITY

18. Security & your data

AmberTrace is built for regulated, high-stakes work. Here is how your data is kept private, where it runs, and the controls you have over it.

Your data stays yours

- **Never shared, never used to train shared models.** Your data, domains, and platforms are visible only within *your* organisation. They are never exposed to other customers, and are never used to train any model used by anyone else.
- **Isolated per organisation.** Every organisation's data is logically isolated, and all access is scoped so that no user can ever see another organisation's work. A **dedicated single-tenant deployment** is available for enterprise customers on request.
- **Encrypted in transit and at rest.** All traffic runs over TLS, and stored data and uploads are encrypted at rest. Organisation administrators can manage and rotate their encryption key on the **Encryption Keys** screen (Section 15).

The model runs on our own infrastructure

The neural component is our own **proprietary model** — built on the open **Llama** family and significantly enhanced by Ambertrace Labs — hosted on **Ambertrace's UK-sovereign infrastructure**. Your data is **never sent to a third-party AI provider** for processing: there is no external LLM in the loop.

Authentication & access controls

- **Multi-factor authentication (MFA).** Every user can enrol for TOTP-based MFA (Section 13, Security tab). Organisation administrators can enforce MFA for all password-login members.
- **Single Sign-On (SSO).** Organisations can configure OIDC-based SSO (Section 13, SSO tab) so members sign in with their corporate identity provider.
- **API key governance.** Keys are scoped (agent vs platform), named, trackable, and instantly revocable (Section 11).

You control sharing — and can see and revoke it

- **No public share-links.** Sharing is strictly within your organisation; there are no link-based or external shares by default. (These can be enabled for enterprise customers on request — see Section 12.)
- **Instant revocation.** Set any domain or platform back to *Private* to withdraw access immediately — affected members are notified.
- **A full audit log.** Every share, access change, and revocation is recorded. Organisation administrators can review the complete history of who was given — or lost — access to what, and when (Section 13, Audit Log tab).

> AVAILABILITY
>
> AmberTrace is offered under an **enterprise licence**. Dedicated deployment, external-sharing options, and bespoke connectors are arranged as part of that licence — talk to us about your

> requirements.
