brandcompete

EnvCraft

User Guide — self-service cloud environments, governed by design

Contents

1 Welcome to EnvCraft 2

2 Getting started 3

3 Your dashboard 6

4 Your preferences 8

5 Projects 11

6 Appliances 14

- Creating an appliance 14
- The appliance workbench 17
- The terminal in the browser 19
- Power schedules 20
- Opening an appliance's web interface 21

7 Templates and the Knowledgebase 22

8 Datasets 26

9 Data sources and scheduled transfers 29

10 Repositories and deploying code 32

11 Costs 34

12 Mason, the assistant 36

13 The audit log 38

14 Administration 39

- Users 39
- Roles & Rights 40
- Vendors 41
- Integrations 43
- Middleware Stacks and Blueprints 44
- Policies & Quotas 45
- Company settings 46

15 Troubleshooting & FAQ 48

A Reference: roles, statuses, glossary 50

1 Welcome to EnvCraft

EnvCraft gives you ready-to-use cloud environments in minutes — without tickets, without infrastructure skills, and within the rules your company has set.

You choose a cloud provider, a software stack and a size. EnvCraft creates the machine, installs and configures the software, mounts your data and tells you how to reach it. Your company decides in advance what is allowed — which clouds, which machine types, how large, how expensive — and EnvCraft checks every request against those rules before anything is created. Machines run in your company's own accounts on Amazon Web Services, Microsoft Azure, Google Cloud or Hetzner Cloud.

The things you work with

Term	What it is
Project	The workspace your work belongs to. Everyone has a personal workspace; a shared project brings colleagues together at Read, Write or Admin. Every appliance, dataset, data source and repository belongs to exactly one project.
Appliance	A cloud machine together with the software configured on it — for example a JupyterLab notebook server, VS Code in the browser, or PostgreSQL with pgAdmin. Every appliance runs Rocky Linux 9.
Template	The recipe for the software half of an appliance, also called a middleware stack. You pick one of 38 when you create an appliance; EnvCraft installs it for you.
Dataset	A folder in managed cloud storage. You fill it from the browser, from a data source or from your appliance, and it appears on your appliances as a normal folder under `/mnt/datasets`.
Data source	A connection to a system your company already runs — a database or a public open-data API. You save queries against it and schedule them to write their results into a dataset.

Alongside them: **repositories** to deploy code from, **power schedules**, a **terminal in the browser**, the **Knowledgebase** of templates and licences, and **Mason**, an assistant you can ask in your own words.

What you can do depends on your role and your project

Every person has a role, and a role is a set of rights. The menu only shows the areas your role may use, and buttons you may not use are greyed out with a hint. Typical roles are **Engineer** (creates and runs appliances and datasets), **Analyst** (works with data shared with them, creates nothing) and **Administrator** (manages the whole company). The full list is in Appendix A.

On top of the role, each project gives its members a level — **Read**, **Write** or **Admin** (chapter 5). An action needs both: the right in your role, and a high enough level in the project. You see your personal workspace, the projects you are a member of, and anything a colleague has shared with you directly. Administrators and auditors can see everything in the company.

NOTE

This guide describes EnvCraft as of 27 September 2026. A few features are visible in the application but not built yet, or are shown as a concept; they are marked `Not built yet` or `Concept` in this guide, and the application says so on screen as well.

2 Getting started

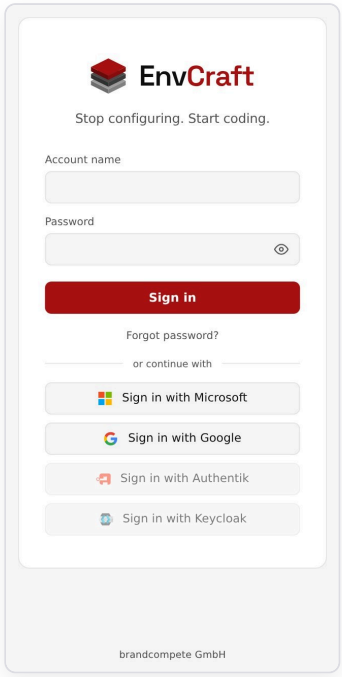
Sign in, find your way around, and prepare the two things every appliance needs.

Signing in

Open the EnvCraft address your company gave you. You can sign in in two ways:

- **With your company account** — click **Sign in with Microsoft** or **Sign in with Google**. EnvCraft only accepts accounts from the directories and email domains your company has allowed.
- **With a password** — enter your **Account name** (your email address) and **Password**, then click **Sign in**. The eye icon shows the password while you type.

Authentik and Keycloak are shown on the page but are not available yet.



The sign-in page

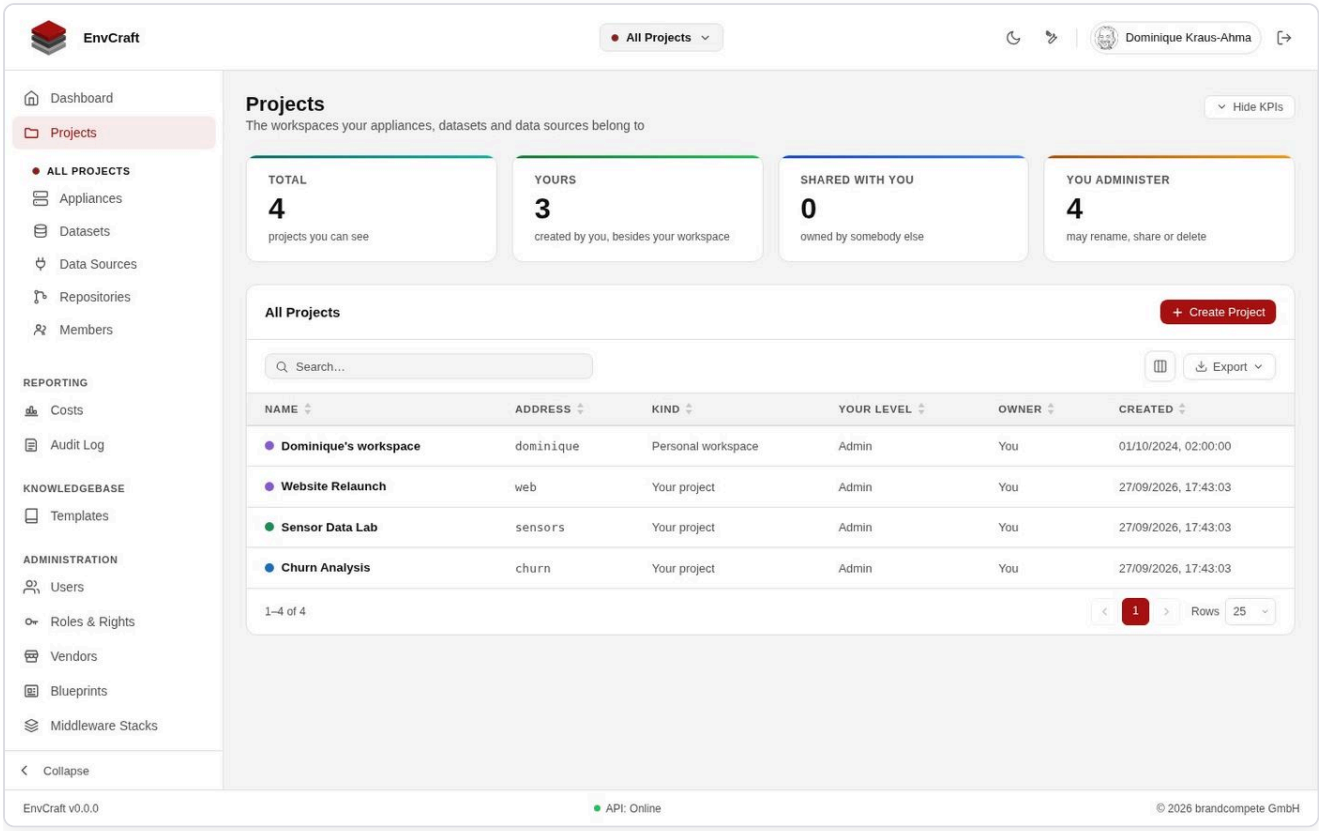
NOTE

The first time you sign in with Microsoft or Google, EnvCraft creates your account with the role **Guest** and a personal workspace. A guest can sign in, read the Knowledgebase and set their own preferences, nothing else — the dashboard says “There is nothing here for you yet.” Ask your administrator to assign you a role.

IMPORTANT

Forgotten your password? Ask your administrator to set a new one. Resetting a password by email is not available yet, even though the sign-in page shows a “Forgot password?” link.

Finding your way around



The application window: sidebar on the left, header with the project switcher at the top, the current page in the middle

- **Sidebar** — the areas of the application: Dashboard and Projects; below them, for the project chosen in the header, Appliances, Datasets, Data Sources, Repositories and Members; **Reporting** (Costs, Audit Log); **Knowledgebase** (Templates); and **Administration** (Users, Roles & Rights, Vendors, Blueprints, Middleware Stacks, Integrations, Policies & Quotas, Settings). **Collapse** at the bottom narrows it to icons.
- **Header** — the **project switcher** (**All Projects** or one project; every list, summary card and create dialog follows it, and EnvCraft remembers your choice), the theme switch (light/dark), the **Tools** button that opens Mason (chapter 12) if your role may use it, and your avatar. Click the avatar to open your preferences; the arrow icon beside it signs you out.
- **Footer** — the EnvCraft version and whether the application can reach its server (**API: Online**).

Working with lists

Every page that lists things — appliances, projects, datasets, data sources, repositories, templates, vendors, integrations, users, the audit log — works the same way:

Element	What it does
Summary cards	Key figures above the list. **Hide KPIs** / **Show KPIs** collapses them; EnvCraft remembers your choice.
Search	Filters the rows as you type. The × clears it.
Filters	Narrows the list by status or kind; a badge shows how many filters are active. **Clear filters** resets them.

Element	What it does
Columns	The column icon opens **Customize columns**: drag to reorder, remove, add, or **Reset to default**.
Sorting	Click a column heading; click again to reverse.
Export	Downloads the current list as Excel (.xlsx), CSV or PDF.
Pages	Browse pages at the bottom; **Rows** sets 10, 25, 50, 100 or All per page.

Clicking a row opens the item's **workbench**: a page with a header naming the item, its most important actions on the right, and tabs underneath.

Before your first appliance

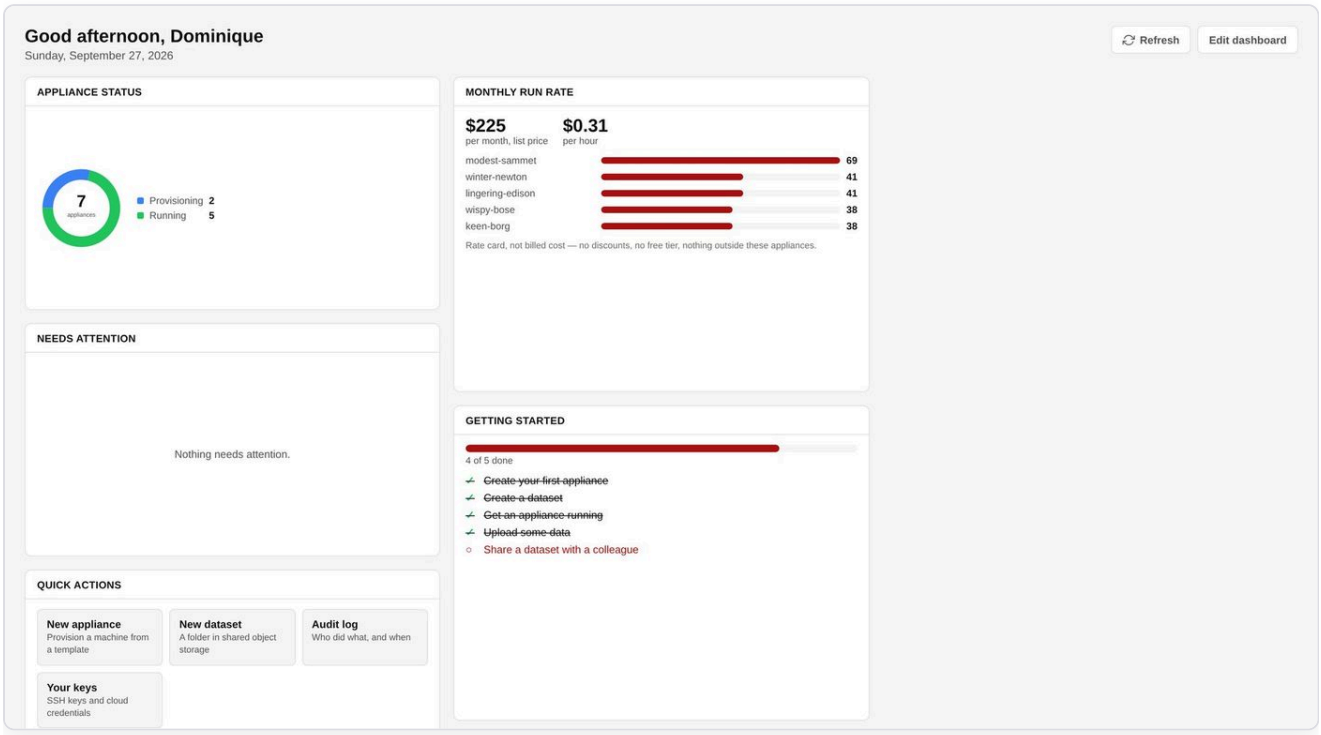
An appliance needs two things from you. Set them up once in your preferences:

1. **An SSH key** — so you can log in to your machines with your own SSH client. Add it under **Linux Identity & SSH Keys**. (The terminal in the browser needs no key of yours.)
2. **A cloud credential** — the cloud account the machine is created in. Usually your company provides a shared one and you need nothing; otherwise add your own under **My Cloud Credentials**.

While you are there, check your **time zone** under Profile & Preferences: power schedules are read in it.

3 Your dashboard

The start page is yours to arrange. Your layout is saved with your account, so it follows you to any browser.

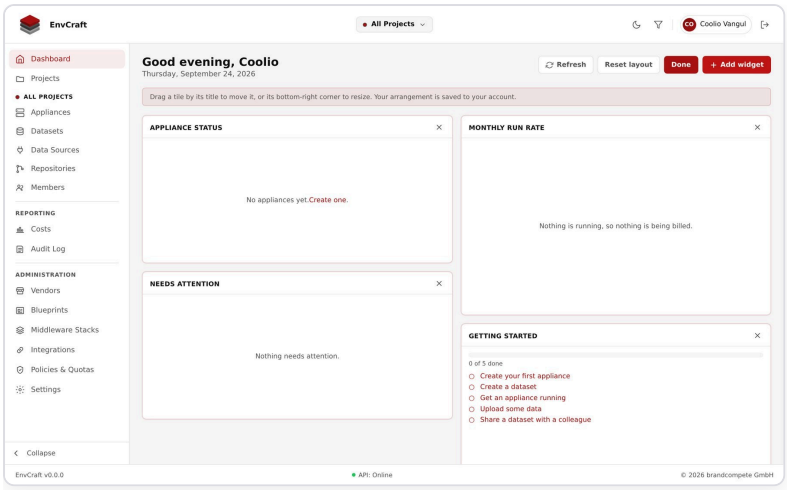


The dashboard: appliance status, monthly run rate, what needs attention, a getting-started checklist and quick actions

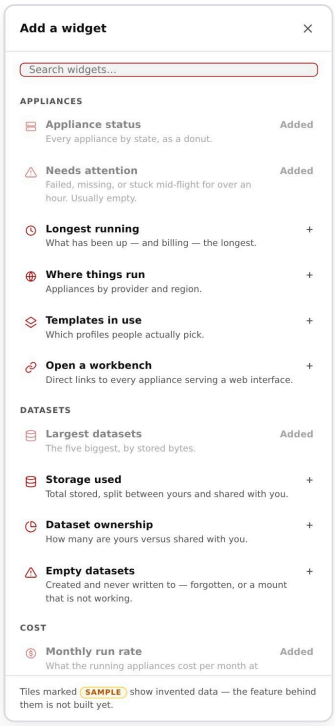
The dashboard is a grid of tiles: the state of your appliances, what needs attention, your monthly run rate, recent activity, a getting-started checklist and quick actions. **Refresh** reloads every tile. Tiles your role cannot use are hidden automatically.

Arranging the tiles

1. Click **Edit dashboard**.
2. Drag a tile by its title to move it; drag its bottom-right corner to resize it; click × to remove it.
3. Click **Add widget** to open the palette, search for a tile and click it to add it.
4. Click **Done**. **Reset layout** returns to the default arrangement.



Edit mode



The widget palette

Available tiles

Group	Tiles
Appliances	Appliance status • Needs attention • Longest running • Where things run • Templates in use • Open a workbench
Datasets	Largest datasets • Storage used • Dataset ownership • Empty datasets
Activity	Recent activity • Activity, 30 days • Refused sign-ins
Cost	Monthly run rate • Cost over time **sample** • Budget **sample**
Operations	Heartbeat • What's new • Messages — all **sample**
Personal	Getting started • Quick actions • Quip of the day **sample** • Did you know **sample**

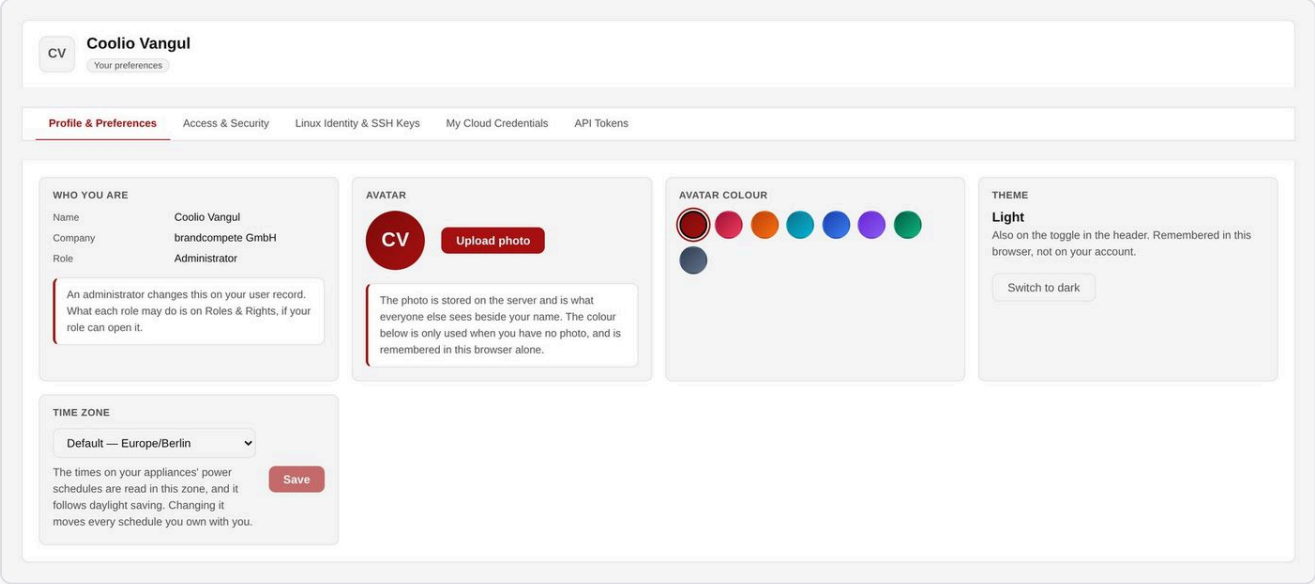
NOT BUILT YET

Tiles marked **sample** show invented data, because the feature behind them is not built yet. The tile, the palette and the tile's footnote all say so.

4 Your preferences

Click your avatar in the header. Everything here concerns you alone.

Profile & Preferences



Profile & Preferences, with the time zone your power schedules follow

- **Avatar** — **Upload photo**, then drag to position it and zoom with the slider or scroll wheel, and click **Apply**. **Remove photo** deletes it.
- **Avatar colour** — the colour of your initials when you have no photo.
- **Theme** — light or dark. The same switch sits in the header.
- **Time zone** — the zone in which the start and stop times of your appliances' power schedules are read. Until you choose one, the installation's default applies (usually Europe/Berlin). Changing it moves every schedule you own with you, and daylight saving is followed automatically.

Access & Security — changing your password

Enter your **Current password**, the **New password** (at least 8 characters, different from the current one) and the new password again, then click **Change password**.

IMPORTANT

Changing your password **signs you out everywhere** — every browser and every session, including this one. Sign in again with the new password. If you sign in with Microsoft or Google, this password does not affect that sign-in.

Linux Identity & SSH Keys



Linux identities: the login names you create appliances with, one of them the default

- **Linux identities** — up to five login names for your appliances: your own, and perhaps one for a shared notebook or a web server. Enter a name (lowercase letters, digits, `_` and `-`, starting with a letter; `root` and EnvCraft's own service name are refused) and click **Add identity**. The **default** is the one the create wizard preselects; you can pick another there each time, or type a one-off name. An appliance keeps the name it was built with, so changing or removing one here never touches a machine that exists.
- **SSH keys** — you can keep up to 4 public keys. Give each a **Name** (for example "Laptop"), paste the **Public key** — the contents of the `.pub` file, never the private key — and click **Add key**. EnvCraft checks the key before saving it.

Accepted key types: `ssh-ed25519` (recommended), `ssh-rsa` with at least 2048 bits, `ecdsa-sha2-nistp256/384/521`, and the hardware-key variants `sk-ssh-ed25519@openssh.com` and `sk-ecdsa-sha2-nistp256@openssh.com`.

> **TIP**
> Don't have a key yet? On Windows, macOS or Linux run `ssh-keygen -t ed25519` in a terminal and paste the contents of the resulting `id_ed25519.pub`.

Removing a key affects only appliances created afterwards; machines that already exist keep the key they were built with.

My Cloud Credentials

A cloud credential is the account an appliance is created in. Your company usually provides one — it appears in the create wizard marked "(company)", and administrators manage it under Vendors (chapter 14). To add your own, click **+ Add credential**, choose the Provider, give it a Label and paste the Credential JSON:

Provider	Credential JSON
Hetzner Cloud	`{"token": "…"}`
Amazon Web Services	`{"accessKey": "AKIA…", "secretKey": "…"}`
Microsoft Azure	`{"tenantId": "…", "clientId": "…", "clientSecret": "…", "subscriptionId": "…", "resourceGroup": ""}` A service principal rather than an access token, which would expire within the hour. `resourceGroup` is optional: set it when the principal may only build into one existing group.
Google Cloud	The service account's JSON key file, pasted whole as Google downloads it (*Keys › Add key*): `{"type": "service_account", "project_id": "…", "private_key": "…", "client_email": "…"}`. The service account needs the roles Compute Admin and Service Account User.

Credentials are encrypted before they are stored and are never shown again. You can edit the label and comment later; to change the secret itself, delete the credential and add it again.

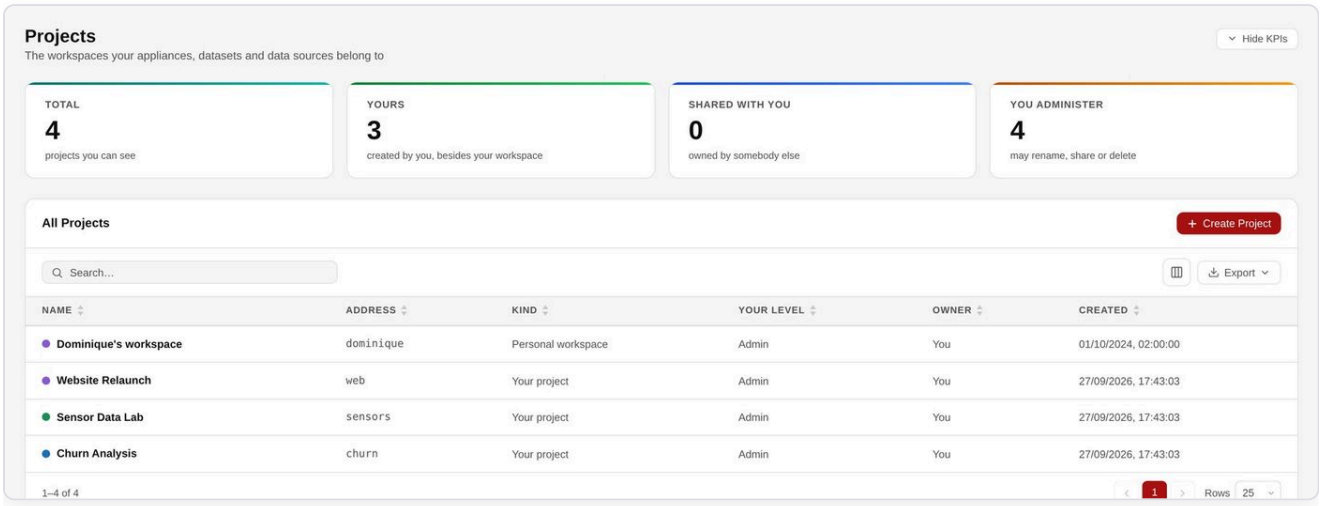
API Tokens

Lists the tokens that let scripts and other applications use EnvCraft on your behalf, and lets you **Revoke** them.

> **NOT BUILT YET**
> Creating new API tokens from this page is not built yet.

5 Projects

A project is where shared work lives: its appliances, datasets, data sources and repositories, the colleagues who may use them, and what it all costs.



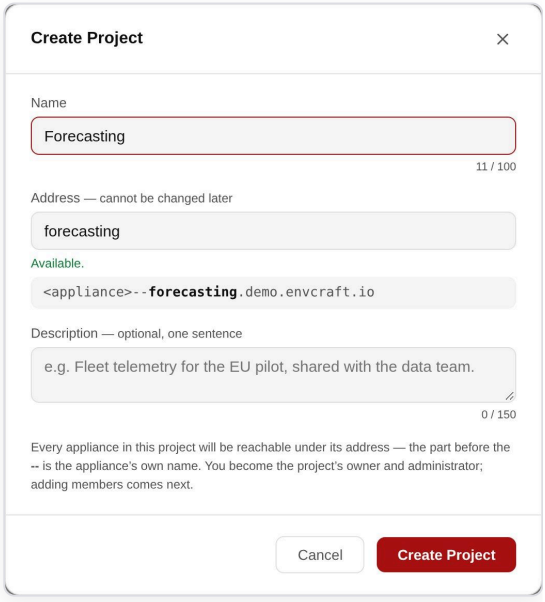
The project list: your personal workspace and the shared projects, with their address, your level and the owner

Your personal workspace

Everyone has one, made the first time they sign in and named after them — for example “Dominique's workspace”. It is yours alone: it can never be shared, and it has no Members tab. Anything you create without choosing a project lands here.

Creating a project

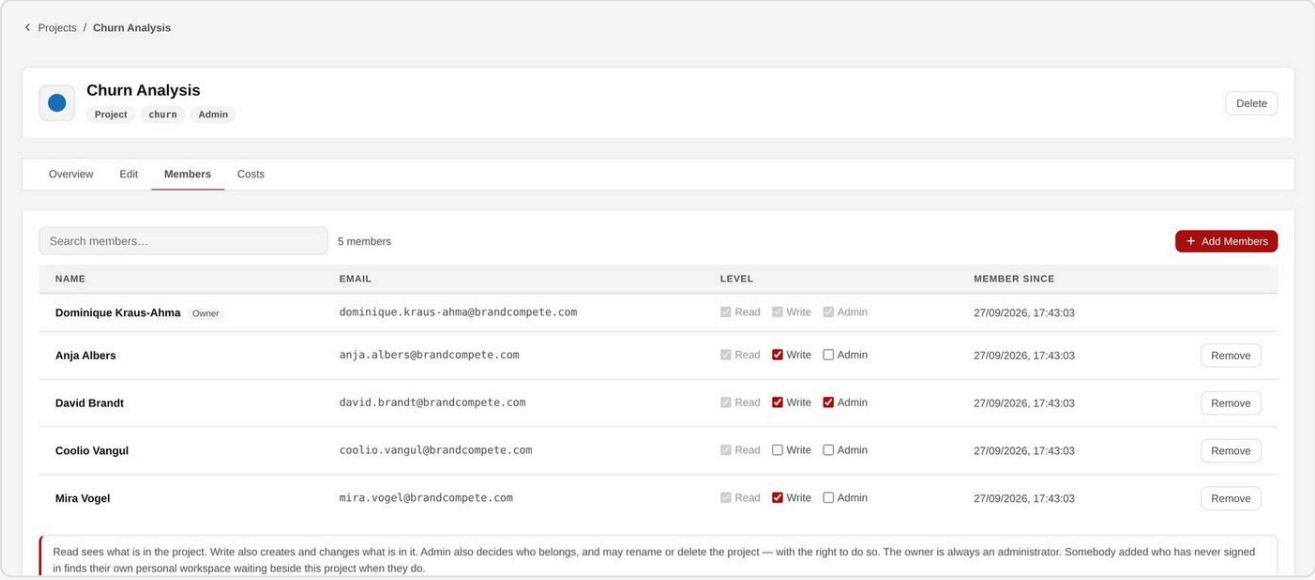
1. On the Projects page, click **Create Project**.
2. Enter a **Name**. EnvCraft suggests an **Address** from it and checks at once whether it is free.
3. Optionally add a one-sentence **Description**.
4. Click **Create Project**. You become the project's owner and administrator; adding members comes next.



IMPORTANT

The **address cannot be changed later**. It is at most 24 characters and becomes part of every appliance address the project publishes. Renaming the project changes only its name.

Members and what they may do



A project's Members tab: each colleague at Read, Write or Admin; the owner is fixed

On the project's **Members** tab, click **Add Members**, search for colleagues, tick a level for each and confirm. Change a level with the tick boxes; **Remove** takes someone out. Added people are notified. The owner cannot be changed or removed.

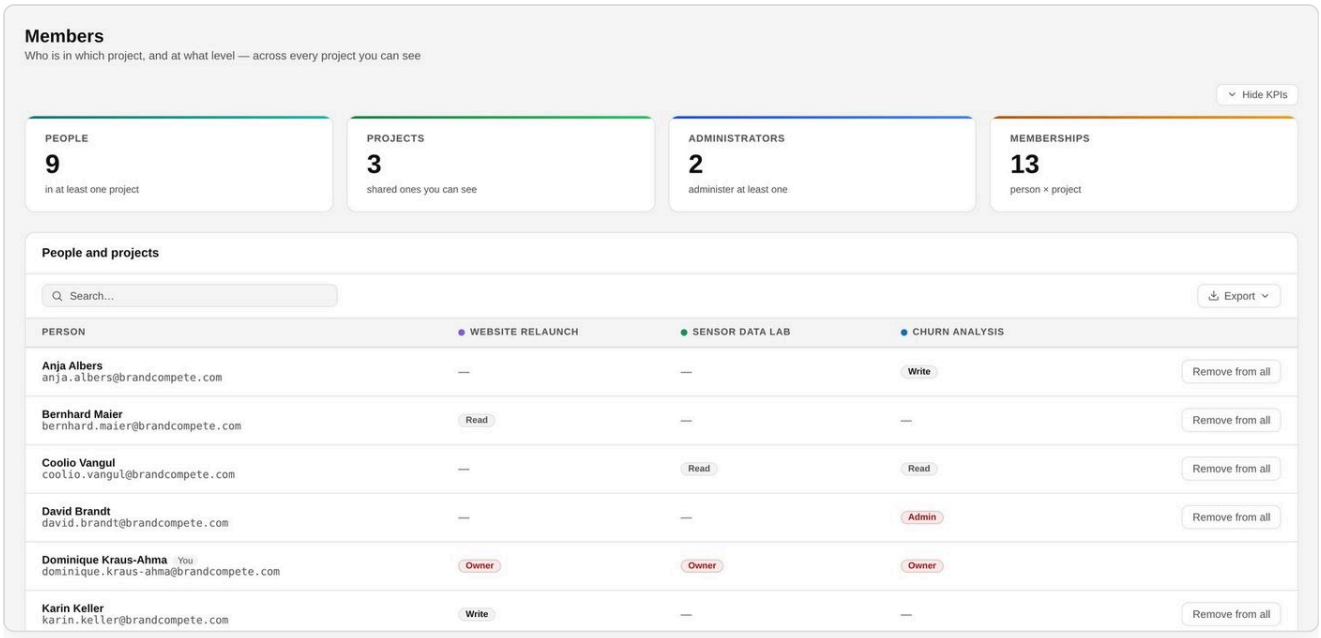
Level	Appliances and data sources	Datasets
Read	See them and their details.	Browse and download the files; mount them read-only on an appliance.
Write	Also use and change them: start, stop, schedule, open the terminal, run queries and transfers.	Also upload and change files; mount them read/write.
Admin	Also destroy or delete them, and manage the members and the project itself.	Also delete them.

The level only opens what your role already allows: an Analyst at Write still creates no appliances. A dataset shared with you directly and through a project gives you the stronger of the two. Removing a member, or lowering their level, reaches their running appliances straight away — the datasets they may no longer use are unmounted.

Project datasets appear on appliances under `/mnt/shared/<owner>/<name>`.

Working in a project

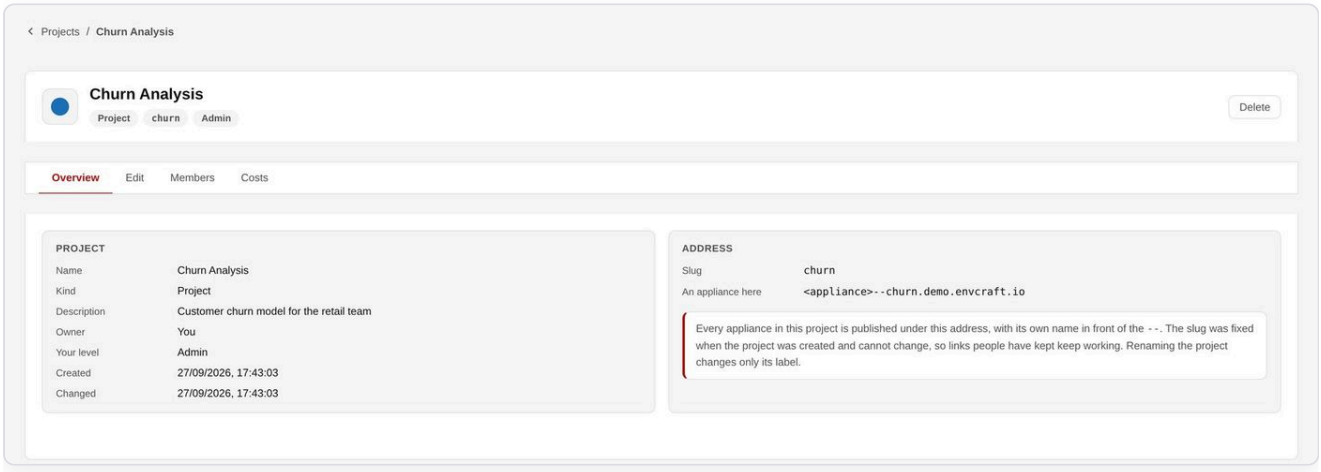
- **The project switcher** in the header narrows every list and summary card to one project, and the create dialogs put new things there. **All Projects** shows everything you can see.
- **Members** in the sidebar lists the members of the chosen project; with **All Projects** it shows a matrix of people and shared projects with their levels, and **Remove from all** takes a person out of every project you administer.
- The project workbench has four tabs: **Overview** (name, owner, your level, the address), **Edit** (name and description), **Members** and **Costs** — the Costs page of chapter 11, narrowed to everything that ever ran in this project, whoever created it.
- A project can be deleted only once it holds nothing any more; a personal workspace cannot be deleted at all.



Members across all projects: who is where, and at which level

Published addresses

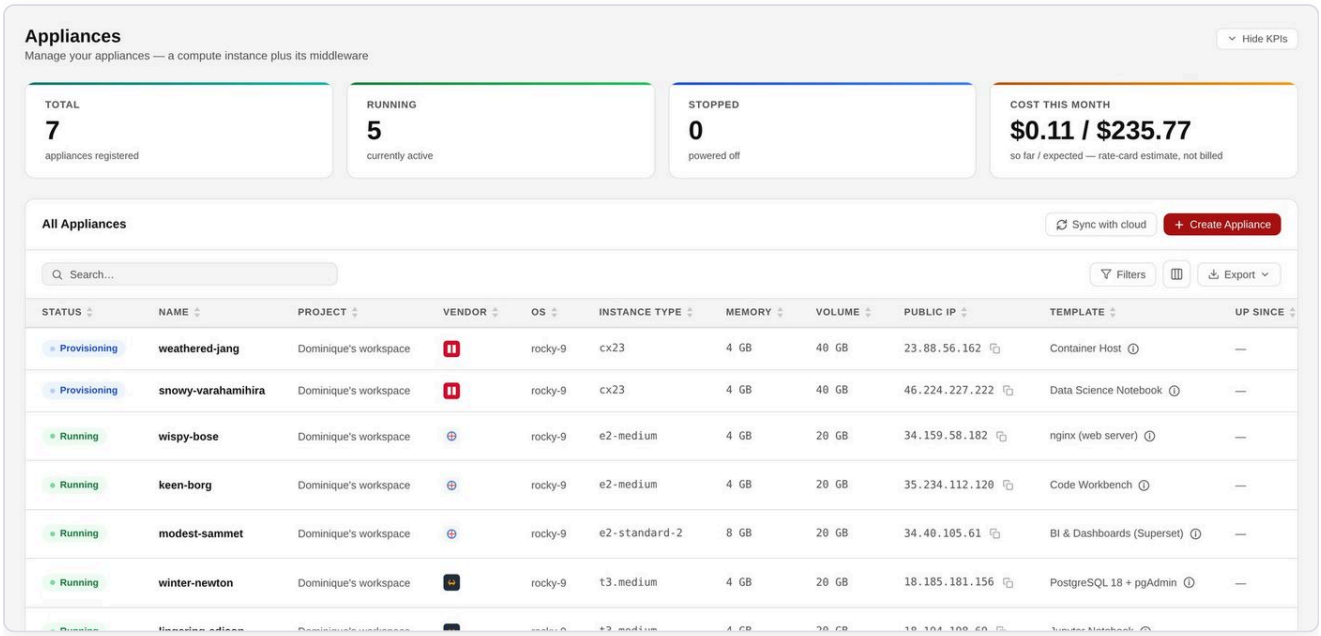
Where your installation has a public address zone, every appliance with a web interface is also published under its project's address: `<appliance>--<project>.<zone>`, for example `keen-borg--churn.envcraft.example.com`, with a proper certificate. The project's Overview shows the pattern. The address is only a name: who may sign in is still decided by the appliance's own sign-in. On installations without a public zone, appliances are opened by their IP address as before.



A project's Overview, with its frozen address and the pattern its appliances are published under

6 Appliances

An appliance is a cloud machine plus the software configured on it. This chapter covers creating, using, connecting to, scheduling and destroying one.



The appliance list, with the project and vendor of each machine

The **Appliances** page lists the appliances in the chosen project with their status, project, vendor, size, public address and template. The summary cards count them and show this month's cost so far and the expected total. The list refreshes itself while anything is starting, stopping or being provisioned.

Actions in the list

- **Sync with cloud** asks each cloud provider what it actually holds and updates the list — useful if something was changed outside EnvCraft.
- **Shutdown** / **Startup** powers a machine off or on. Shutting down disconnects anyone working on it and loses unsaved work, but the disk, the address and the datasets survive.
- **Open** (when the appliance serves a web interface and is running) opens it in a new browser tab.
- The **terminal icon** opens a shell on the appliance in a tab of its own (see *The terminal in the browser*).
- **View log** (when the status is **Error**) shows what went wrong.

Destroying an appliance is deliberately not in the list — it is in the workbench, where you can see what you are about to delete.

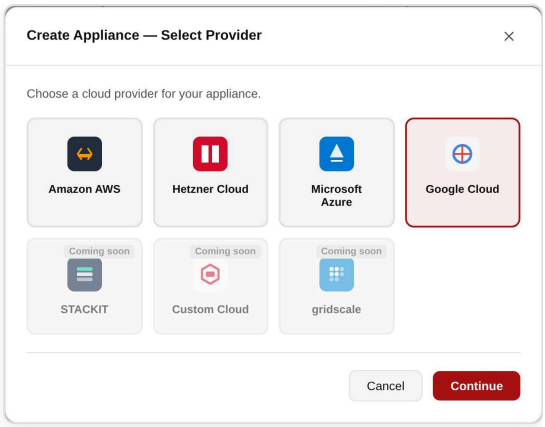
Creating an appliance

Click **Create Appliance**. The wizard has four steps. You can also ask Mason to prepare one for you (chapter 12).

Step 1 — Select provider

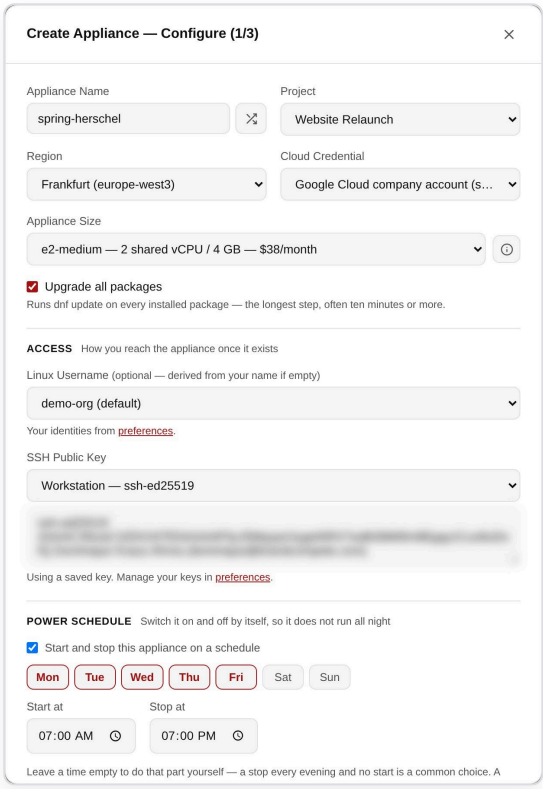
Choose where the machine will run: **Amazon AWS**, **Hetzner Cloud**, **Microsoft Azure** or **Google Cloud**. STACKIT, custom clouds and gridscale are marked “Coming soon”.

Only the vendors your company has switched on can be chosen; a vendor it has switched off is greyed out.



Step 2 — Configure

- **Appliance Name** — a suggestion is filled in; the shuffle button offers another. You can rename the appliance later.
- **Project** — where the appliance belongs; preset from the project switcher.
- **Region** and **Cloud Credential** — your own credential or your company's (marked “(company)”). The dropdowns say what they are loading while they ask the vendor.
- **Appliance Size** — live from the vendor: what it actually sells in that region, less any types your company has withheld, with the monthly price. Long lists are grouped by family; the info icon shows processor, memory, disk and price.
- **Upgrade all packages** — runs a full operating-system update while the machine is built: the longest step, often ten minutes or more. Untick it for a quick demonstration machine. Your company may enforce it, and below 4 GB of memory it is skipped either way.
- **Linux Username** — your default Linux identity, another of yours, or **Other…** for a one-off name.
- **SSH Public Key** — a saved key, or a pasted one (your company may require a saved key).
- **Power schedule** — optionally, weekdays and a start and stop time. Your company may preselect one.



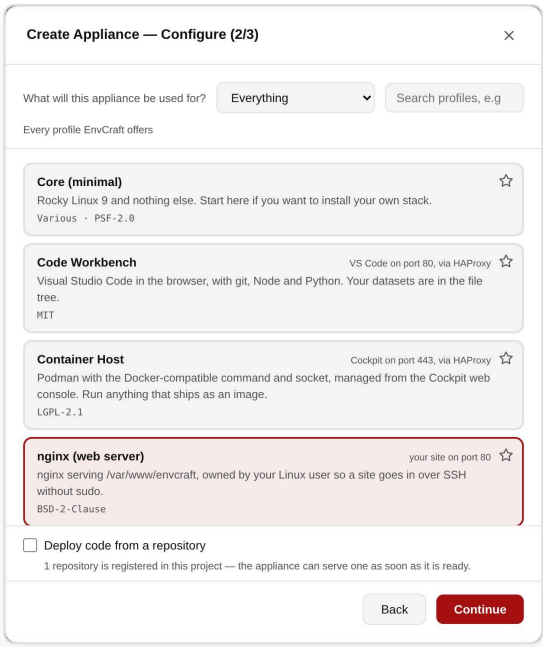
NOTE

Hetzner prices are published in euros; EnvCraft converts them to US dollars at the day's ECB reference rate so that every figure, budget and chart uses one currency. All prices are net list prices and include the primary IPv4 address.

Step 3 — Pick the software

Choose one of the 38 templates (chapter 7). **What will this appliance be used for?** groups them — My favourites, Everything, Developer, Data scientist, DevOps and Citizen developer — and the search box finds a template across all groups. Each card says which web interface it publishes and lists its licences; the star adds it to your favourites.

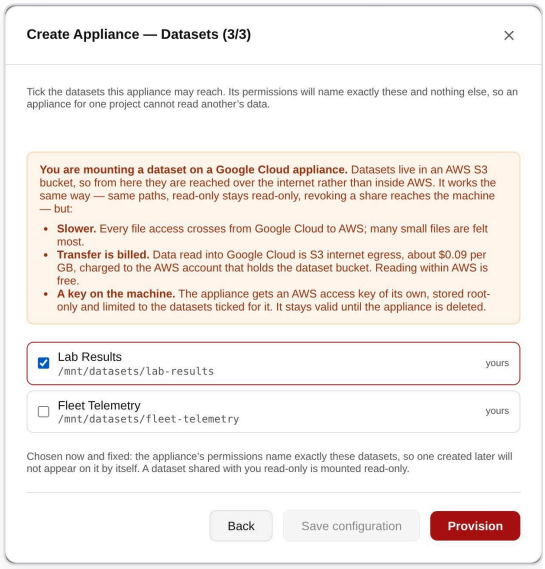
For the **nginx web server**, tick **Deploy code from a repository** to have the machine serve a registered repository from the start (chapter 10). The **Build Runner** asks for the repository it should serve.



Step 4 — Datasets, then Provision

Tick the datasets this appliance may use. Your own appear at `/mnt/datasets/<name>`; datasets shared with you or through a project appear under `/mnt/shared/…`, read-only or read/write as granted.

Click **Provision**. EnvCraft checks your company's limits first; if a limit would be exceeded, it tells you which one and creates nothing.



IMPORTANT

Datasets outside AWS. Datasets live in AWS storage. An appliance on Azure, Google Cloud or Hetzner reaches them over the internet, and the wizard says so before you tick anything: access is **slower** (many small files are felt most), data read into the other cloud is **billed as AWS transfer** (about $0.09 per GB, to the account that holds the datasets), and the appliance gets **an access key of its own**, stored root-only and limited to the ticked datasets. Paths, read-only access and revoking a share behave exactly as on AWS.

NOT BUILT YET

Save configuration — saving a setup to reuse later — is not built yet.

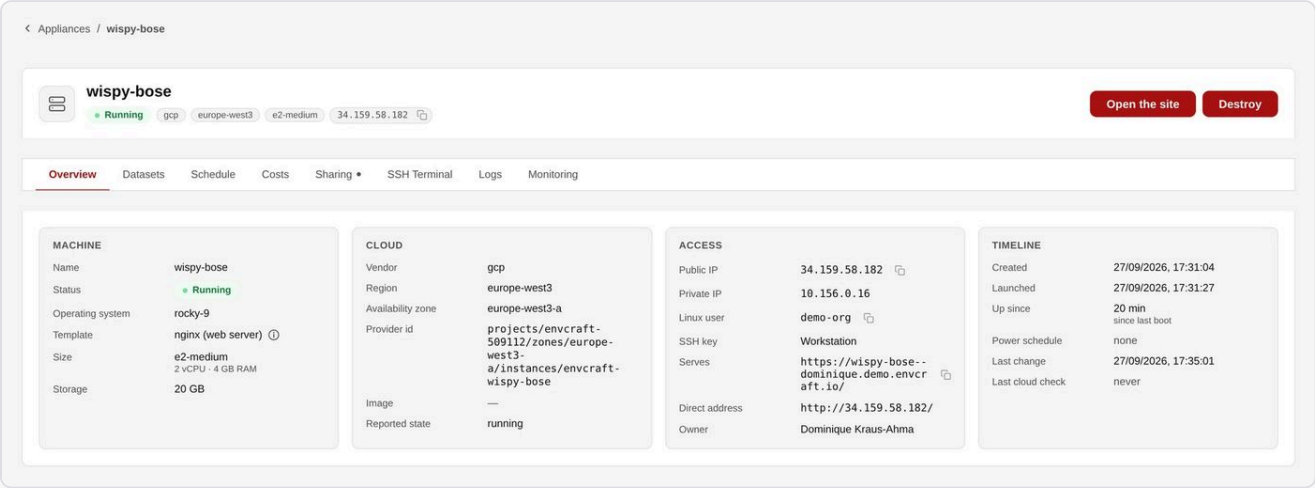
Watching the provisioning

After you click Provision, the wizard shows a live log: EnvCraft first creates the machine in the cloud, then installs and configures the software. You can click **Close — keep provisioning** at any time; the work continues and the log remains on the appliance's **Logs** tab. The appliance shows **Provisioning** until it is **Running**.

Every appliance is set up the same way underneath: daily security updates (never with an automatic reboot), fail2ban protecting SSH, a set of administration tools, and a sign-in banner that shows the owner, the published address, load, memory and some useful commands. Only SSH and the web ports 80 and 443 are open to the outside.

The appliance workbench

Click an appliance in the list to open its workbench.



The Overview tab — here a web server on Google Cloud

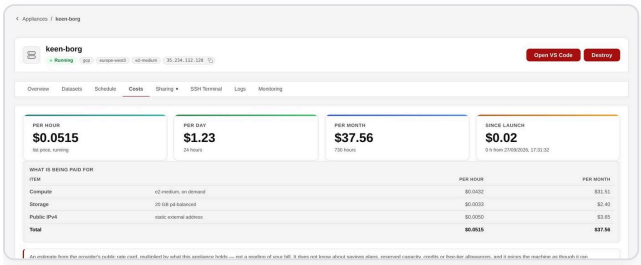
The header

- The name with a pencil icon to rename it. Renaming changes only the label you see; the machine keeps the name it was created with.
- Status, vendor, region and size, and the public IP address with a copy button.
- **Open …** (for example “Open VS Code” or “Open the site”) when the appliance serves a web interface.
- **Destroy** — see below.

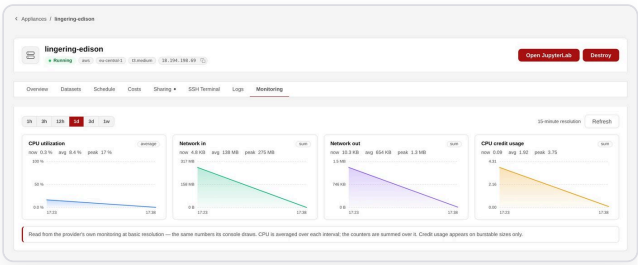
The tabs

Tab	What you find there
Overview	Four panels. **Machine**: name, status, operating system, template, size, storage. **Cloud**: vendor, region, zone, provider id, image, reported state. **Access**: public and private IP, the Linux user with **Copy ssh command**, the SSH key, the web addresses and their sign-in details, the owner. **Timeline**: created, launched, up since, power schedule, last change, last cloud check.
Datasets	Change which datasets are mounted. Tick or untick, then **Apply to appliance**; EnvCraft reconfigures the machine.
Schedule	The appliance's power schedule — see *Power schedules* below.
Costs	Estimated cost per hour, day and month, cost since launch, and an itemised table of what is being paid for.
Sharing	**Not built yet** — granting a single colleague access to one appliance. Share through a project instead (chapter 5).
SSH Terminal	A shell on the machine in the browser — see below.

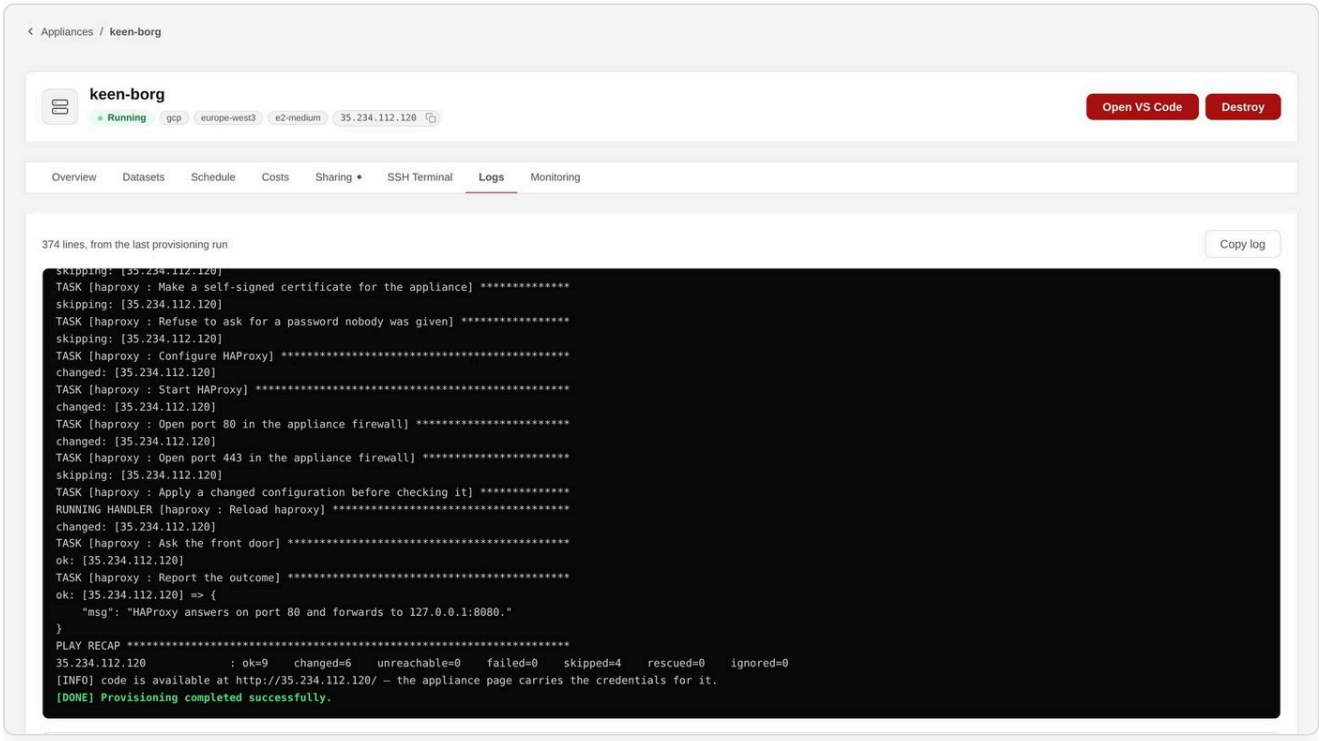
Tab	What you find there
Logs	The provisioning log — live while it runs, kept afterwards. **Copy log** copies it, for example to send to support.
Monitoring	CPU utilisation, network in and out, and CPU credit usage, over the last hour up to the last week. Read from the provider's own monitoring — available for AWS appliances.



The Costs tab



The Monitoring tab of an AWS appliance



The Logs tab: every provisioning step, ending in “Provisioning completed successfully”

Connecting with your own SSH client

On the Overview tab, click **Copy ssh command** next to the Linux user and paste it into a terminal. It connects with the SSH key you chose when creating the appliance. Your datasets are under `/mnt/datasets` and `/mnt/shared`.

NOTE

Directory listings of a mounted dataset are cached on the machine for about 30 seconds, so a file uploaded in the browser can take a moment to appear in a shell.

Destroying an appliance

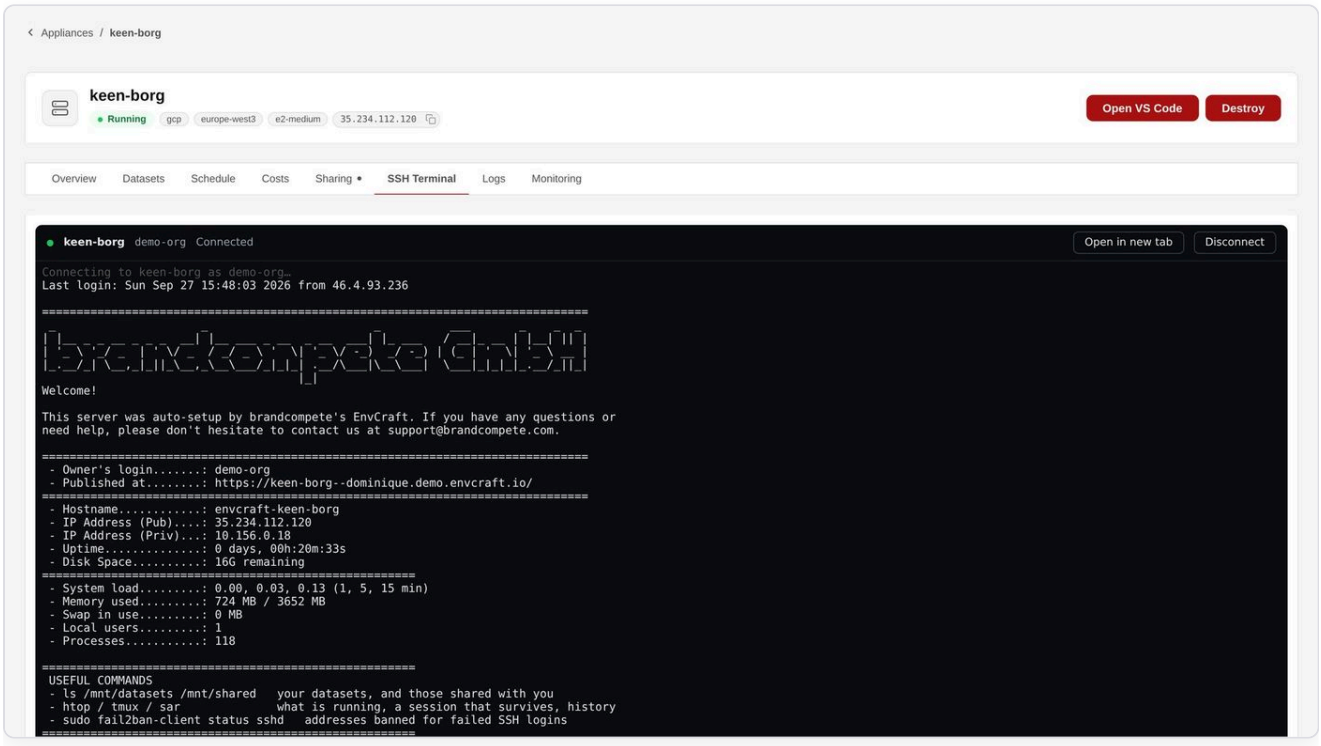
Click **Destroy** in the workbench header and confirm with **Destroy appliance**. The row is marked as deleting at once, and the workbench returns to the list.

> **IMPORTANT**
> Destroying deletes the machine and everything stored on it — the disk, its public address and its cloud permissions. It cannot be undone. Files in mounted datasets are not affected: they live in the dataset, not on the machine. In a project, destroying needs Admin or ownership.

If an appliance shows **Not in cloud** — it was deleted outside EnvCraft — the button reads **Remove** and only removes the entry from EnvCraft.

The terminal in the browser

The **SSH Terminal** tab opens a shell on the appliance, as its Linux user, without an SSH client or a key of your own. **Open in new tab** moves it to a browser tab of its own; the terminal icon in the appliance list does the same. **Disconnect** ends the session.



A session in the SSH Terminal tab, with the appliance's sign-in banner

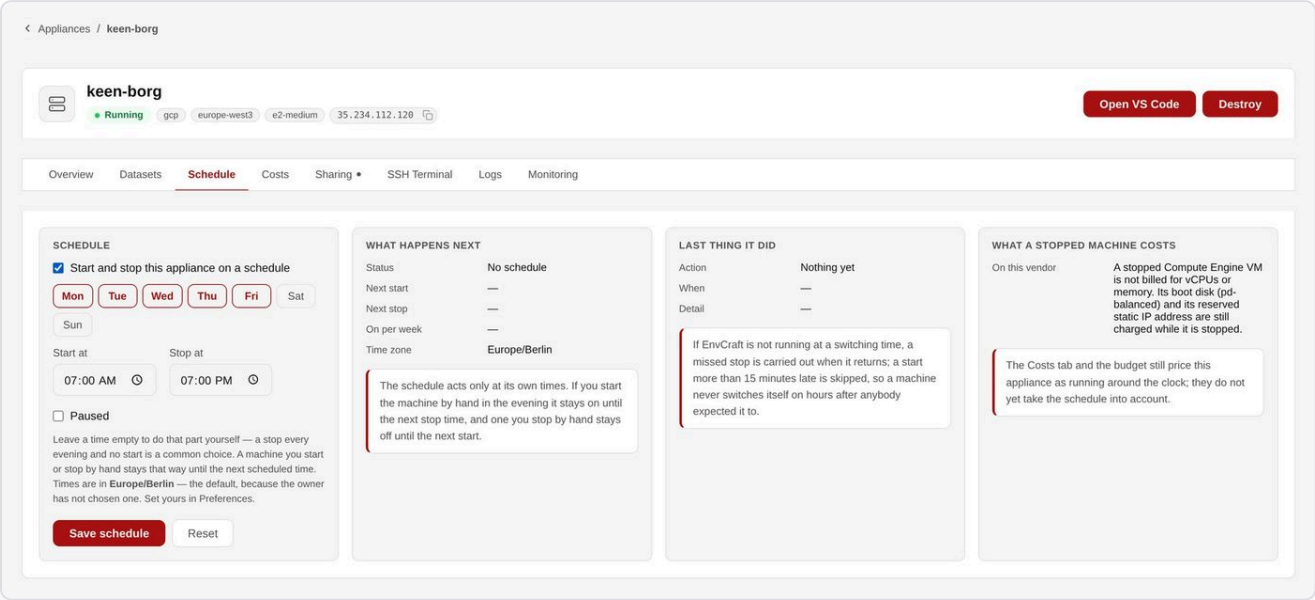
- **Who may open one:** the owner, or a project member with Write — and in both cases the right to connect to appliances in your role. The appliance must be running.
- **How it is secured:** EnvCraft signs every session with an SSH certificate that is valid for five minutes, names you, and allows an interactive shell only — no forwarding. There is no key to copy around and nothing to revoke afterwards.
- **Limits:** a session ends when you close the tab, after 30 minutes without input, or after 8 hours. You can have up to five sessions at a time, and an appliance takes up to five.
- **Audit:** opening and closing a session are recorded, with its duration. What you type is not recorded.

NOTE

Appliances built before the terminal existed do not trust EnvCraft's certificates; their terminal icon is grey with the hint “rebuild the appliance to use it”. Mason and the MCP server cannot open a terminal.

Power schedules

A development machine rarely needs to run all night. A power schedule starts and stops an appliance by itself: on chosen weekdays, at a start time and a stop time. Set one in step 2 of the wizard, or later on the workbench's **Schedule** tab.



The Schedule tab: weekdays, start and stop time, what happens next and what a stopped machine still costs

1. Tick **Start and stop this appliance on a schedule**.
2. Choose the weekdays and the **Start at** and **Stop at** times. Leave a time empty to do that part yourself — a stop every evening and no start is a common choice.
3. Click **Save schedule**. **Paused** suspends a schedule without deleting it; **Remove schedule** deletes it.

- **A schedule acts only at its own times.** If you start the machine by hand in the evening, it stays on until the next stop time; one you stop by hand stays off until the next start.
- **Times are in the owner's time zone**, set in their preferences.
- If EnvCraft itself is not running at a switching time, a missed stop is carried out when it returns; a start more than 15 minutes late is skipped.

NOTE

Power schedules work on **AWS, Azure and Google Cloud**. A stopped AWS machine costs only its storage and address; a stopped Azure machine is deallocated, and a stopped Google Cloud machine is not billed for processor and memory — in both cases the disk and static address still are. EnvCraft cannot power Hetzner machines on a schedule — and Hetzner bills a powered-off server in full, so a schedule there would save nothing. The Costs tab and the budget still price a scheduled appliance as if it ran around the clock.

Opening an appliance's web interface

31 of the 38 templates come with a web interface. Click **Open** in the list or the workbench header. Every web interface is reached on the machine's ordinary web port — applications that listen elsewhere sit behind a small local front door (HAProxy) on the machine, so links carry no port number. How you sign in depends on the template, and the Access panel shows what applies:

Sign-in	Templates	How you get in
At the front door	Jupyter, Data Science, R Workbench, R + Shiny, scikit-learn, TensorFlow, PyTorch, Model Serving, Experiment Tracker, Spark, Streamlit, Chainlit, Dash, Workflow Automation, Observability, SQL Server and Oracle (CloudBeaver)	Your browser asks for a user name and password: the Linux user and the password under Overview › Access.
The application's own	Code Workbench (password only), pgAdmin, Cockpit, Superset, the RabbitMQ console, the Traefik dashboard, Baserow, BookStack, MediaWiki, Appsmith, Airflow	The account and password under Overview › Access. EnvCraft creates the first account for you, so nobody else can claim it.
None — public	nginx, Node Workspace, Flask Web App	Your own site or app on port 80.

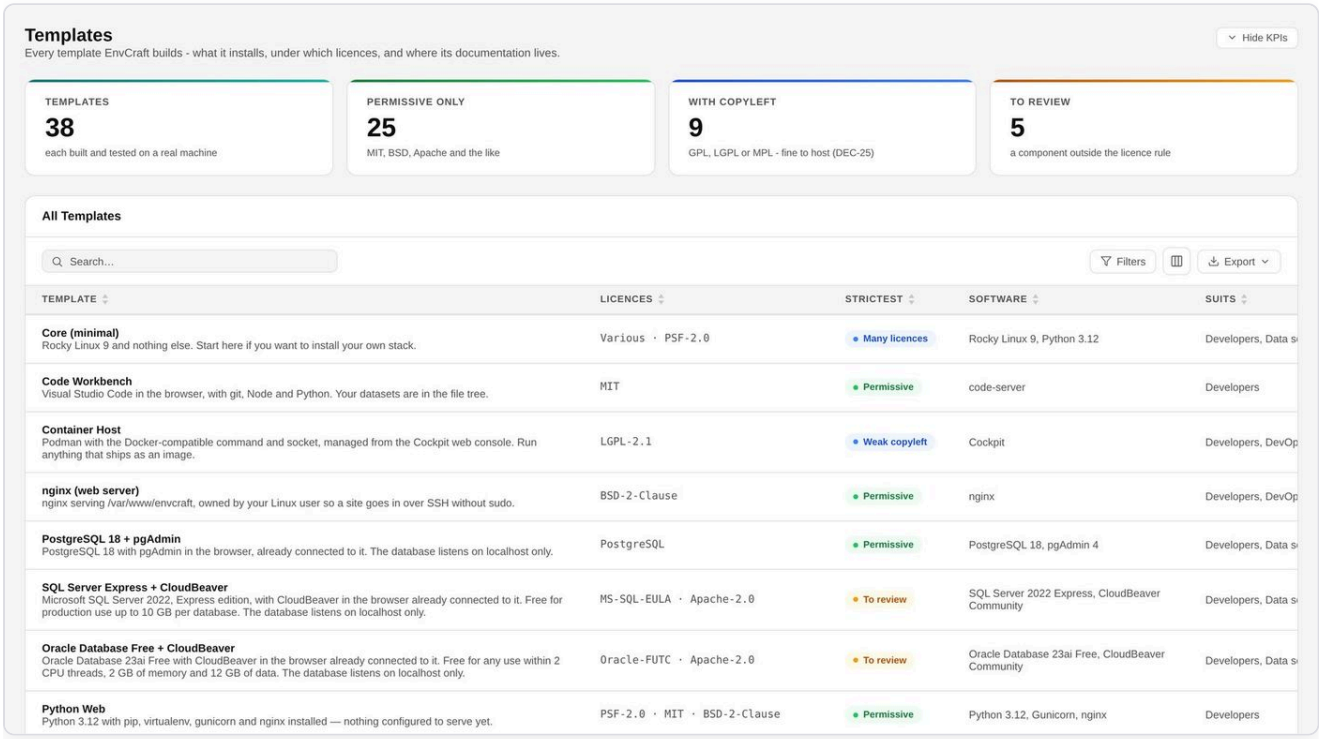
Credentials are generated per appliance and have copy buttons. Cockpit on the Container Host is served over HTTPS with a self-signed certificate, so your browser warns once — accept it to continue. Templates without a web interface (Core, Python Web, PostgreSQL 16, JVM Workspace, Analytics Warehouse, DuckDB, Build Runner) are used over SSH.

7 Templates and the Knowledgebase

38 templates, each built and tested on a real machine at its minimum size — and a Knowledgebase that says what each installs and under which licence.

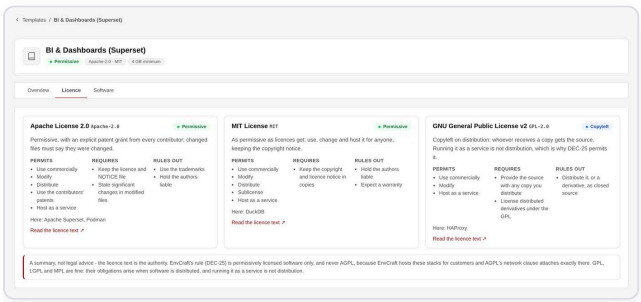
The Knowledgebase

Knowledgebase › Templates in the sidebar is open to everybody, guests included. It lists every template with its licences, the strictest kind of licence among them (a coloured pill: permissive, weak copyleft, copyleft, or **To review** in amber), the software it installs, whom it suits and its minimum memory. Filter by role, licence or kind, or search.

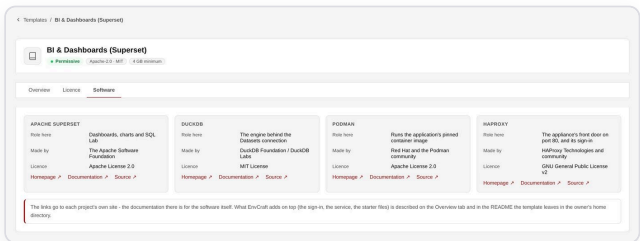


The Knowledgebase: 38 templates, their licences and the strictest one

Click a template for its own page with three tabs: **Overview**; **Licence** — one card per licence with a plain summary of what it permits, what it requires and what it rules out, and a link to the full text; and **Software** — every component with its role, maker, licence and links to homepage, documentation and source code.



The Licence tab



The Software tab

NOTE

The summaries are a guide, not legal advice; the licence text is the authority. Nothing is offered to your users over the network under the AGPL, except in one template still under review (R + Shiny, whose Shiny Server is AGPL); PM2 in the Node Workspace is AGPL too, but only supervises your app and is never offered over the network. Three components are proprietary but free of charge under their vendors' terms: SQL Server Express, Oracle Database Free and the Netdata dashboard.

The catalogue

Template	What it installs	Web interface	Licences	RAM
Core (minimal)	Rocky Linux 9 and Python 3.12, nothing else — start here to install your own software.	—	various, PSF	1 GB
Code Workbench	VS Code in the browser (code-server), git, Node.js and Python. Your datasets are in the file tree.	VS Code, password	MIT	2 GB
Container Host	Podman with Docker-compatible commands and socket, managed from the Cockpit console.	Cockpit, port 443	LGPL-2.1	2 GB
nginx (web server)	nginx serving `/var/www/envcraft`, owned by your Linux user; can be filled from a repository.	your site, public	BSD-2	1 GB
PostgreSQL 18 + pgAdmin	PostgreSQL 18 (listening on the machine only) and pgAdmin, already connected.	pgAdmin	PostgreSQL	2 GB
SQL Server Express + CloudBeaver	SQL Server 2022 Express and CloudBeaver in the browser, already connected. Express limits: 10 GB per database, up to 4 cores, about 1.4 GB buffer memory.	CloudBeaver, front door	MS SQL EULA, Apache-2.0 To review	4 GB
Oracle Database Free + CloudBeaver	Oracle Database 23ai Free and CloudBeaver, already connected. Free limits: 2 CPU threads, 2 GB memory, 12 GB of user data.	CloudBeaver, front door	Oracle FUTC, Apache-2.0 To review	4 GB
Python Web	Python 3.12 with pip, virtualenv, gunicorn and nginx — installed, not configured.	—	PSF, MIT, BSD-2	2 GB
PostgreSQL 16	The database server only. Reachable over SSH.	—	PostgreSQL	2 GB
Node Workspace	Node.js 22 with pnpm and PM2, nginx on port 80 in front of your app.	your app, public	MIT, AGPL (PM2), BSD-2 To review	2 GB
Flask Web App	Flask with gunicorn behind nginx.	your app, public	BSD-3, MIT, BSD-2	2 GB
JVM Workspace	OpenJDK 25 with Maven and Gradle 9 — a build box reached over SSH.	—	GPL-2.0 + CE, Apache-2.0	4 GB
Jupyter Notebook	JupyterLab with pandas, numpy and matplotlib; opens on a notebook listing your datasets.	JupyterLab, front door	BSD-3	2 GB
Data Science Notebook	JupyterLab with polars, DuckDB, scikit-learn and seaborn; DuckDB queries files in `/mnt/datasets` in place.	JupyterLab, front door	BSD-3, MIT	4 GB
R Workbench	R in JupyterLab (IRkernel) with the tidyverse. The first build is slow.	JupyterLab, front door	GPL-2.0, MIT, BSD-3	4 GB
R + Shiny	R and Shiny Server: put an app in `/srv/shiny-server` and it is live.	Shiny, front door	GPL-2.0, AGPL (Shiny Server) To review	4 GB
Experiment Tracker	MLflow tracking, with artifacts on a dataset where one is mounted.	MLflow, front door	Apache-2.0	4 GB

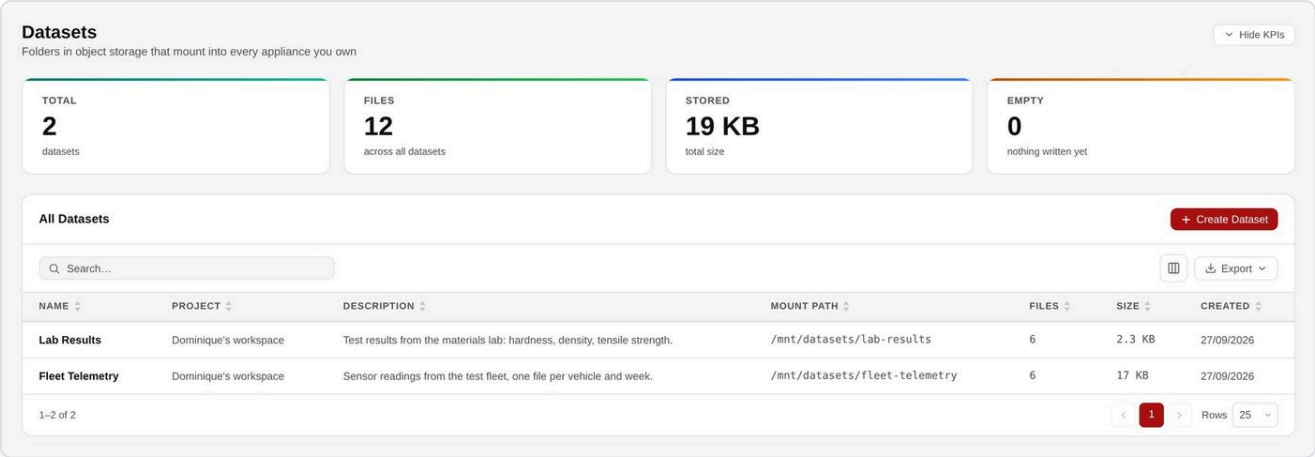
Template	What it installs	Web interface	Licences	RAM
Analytics Warehouse	DuckDB with dbt and a starter project over `/mnt/datasets`.	—	`MIT, Apache-2.0`	4 GB
Model Serving	FastAPI predictions for the joblib models in `~/models`.	API docs, front door	`MIT, BSD-3`	2 GB
Chat Apps (Chainlit)	Chainlit with the Anthropic SDK — a chat app that uses Claude once you set a key.	Chainlit, front door	`Apache-2.0, MIT`	2 GB
Machine Learning (scikit-learn)	JupyterLab with scikit-learn, XGBoost and LightGBM.	JupyterLab, front door	`BSD-3, Apache-2.0, MIT`	4 GB
Deep Learning (TensorFlow)	JupyterLab with TensorFlow and Keras, **CPU only**.	JupyterLab, front door	`Apache-2.0, BSD-3`	8 GB
Deep Learning (PyTorch)	JupyterLab with PyTorch and torchvision, **CPU only**.	JupyterLab, front door	`BSD-3`	8 GB
Apache Spark	A single-node Spark 4 with PySpark.	master UI, front door	`Apache-2.0, GPL-2.0 + CE`	8 GB
DuckDB	The DuckDB command line and a Python environment over `/mnt/datasets`.	—	`MIT, BSD-3`	2 GB
Workflow Orchestration (Airflow)	Apache Airflow 3 with its own PostgreSQL.	Airflow, own sign-in	`Apache-2.0, PostgreSQL`	4 GB
BI & Dashboards (Superset)	Apache Superset, with DuckDB over the mounted datasets.	Superset, own sign-in	`Apache-2.0, MIT`	4 GB
Data Apps (Streamlit)	Streamlit with DuckDB.	Streamlit, front door	`Apache-2.0, MIT`	2 GB
Dashboards in Python (Dash)	Dash and Plotly under gunicorn.	Dash, front door	`MIT`	2 GB
Build Runner	A GitHub Actions runner registered with a repository of the project, removed again when the appliance is destroyed.	—	`MIT`	4 GB
Observability	Netdata with its dashboard, and Prometheus scraping it.	Netdata, front door	`GPL-3.0, Netdata licence, Apache-2.0` **To review**	4 GB
Edge & Ingress	Traefik on ports 80 and 443 with file routes and Let's Encrypt.	dashboard, own sign-in	`MIT`	1 GB
Cache & Queue	Valkey and RabbitMQ, listening on the machine; the RabbitMQ console published.	RabbitMQ, own sign-in	`BSD-3, MPL-2.0`	2 GB
Workflow Automation	Node-RED.	Node-RED, front door	`Apache-2.0`	2 GB
Data Tables	Baserow, with your account created first and sign-ups closed.	Baserow, own sign-in	`MIT`	4 GB
Internal App Builder	Appsmith Community, you as administrator, sign-ups closed.	Appsmith, own sign-in	`Apache-2.0`	4 GB
Docs & Knowledge	BookStack on MariaDB, with your account instead of the default one.	BookStack, own sign-in	`MIT, GPL-2.0`	2 GB
Wiki (MediaWiki)	MediaWiki on PostgreSQL, your Linux user as administrator.	MediaWiki, own sign-in	`GPL-2.0, PostgreSQL`	2 GB

Deep learning runs on the processor until appliances can be given a GPU, and Spark is a single machine rather than a cluster. If you outgrow SQL Server Express or Oracle Database Free, your company licenses a paid edition itself. The Oracle image is a widely used community build (gvenzl/oracle-free) rather than one published by Oracle.

The **Middleware Stacks** page under Administration shows the same catalogue with how many appliances run each template; a template's **Provisioning** tab explains, step by step, what EnvCraft does when you create an appliance with it.

8 Datasets

A dataset is a folder in managed cloud storage that appears on your appliances as an ordinary folder — on every cloud.



The dataset list, with the project, mount path, file count and size of each

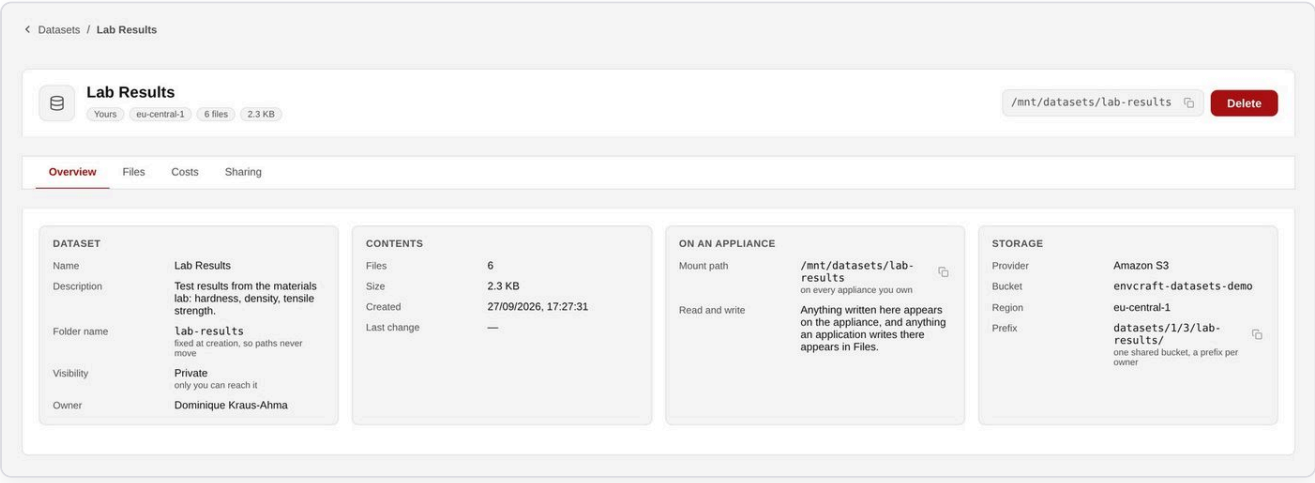
Creating a dataset

1. On the Datasets page, click **Create Dataset**.
2. Enter a **Name** (for example “Churn 2026”), optionally a **Description**, and the project it belongs to.
3. Check the preview **Folder on your appliances** — EnvCraft turns the name into a folder name such as `/mnt/datasets/churn-2026`.
4. Click **Create**.

NOTE

The folder name is fixed when the dataset is created, so paths in your scripts never break. Renaming the dataset later only changes its label. Your company may add a prefix or suffix to every folder name.

The dataset workbench



Overview: the dataset, its contents, where it appears on an appliance, and where it is stored

Tab	What you find there
Overview	Name, description, folder name, visibility and owner; file count and size; the mount path on your appliances; the storage location.
Files	Browse, upload, download and delete files. **Refresh** reads the folder again; the time of the last listing is shown.
Costs	The estimated storage cost of the files it holds.
Sharing	Who else may use the dataset.

Uploading and managing files

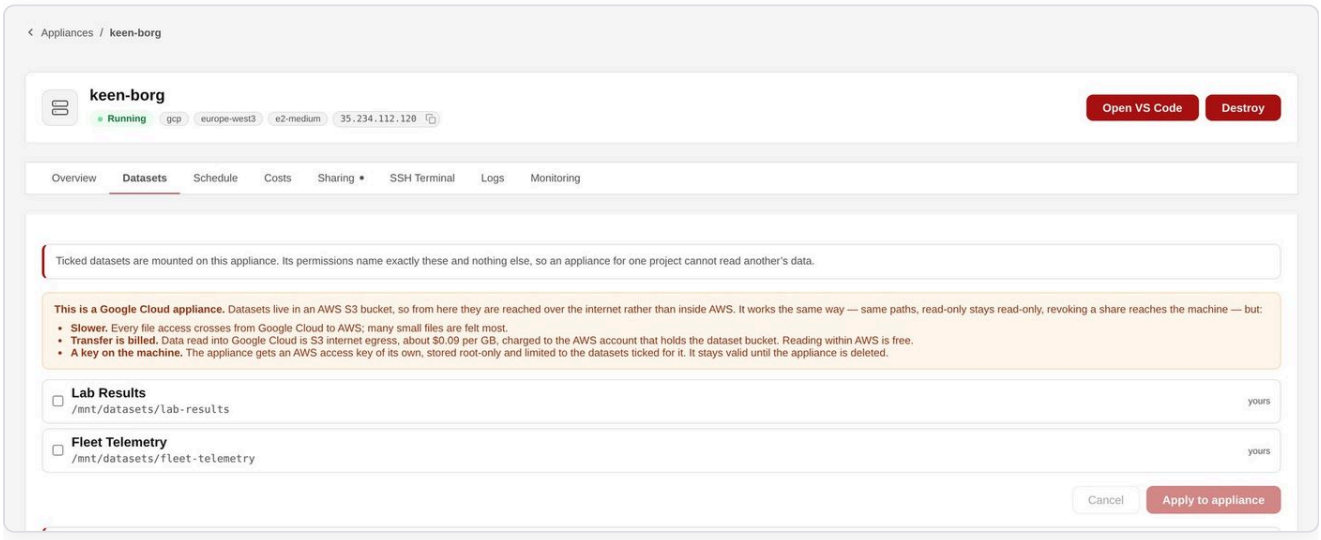
On the Files tab, click **Upload Files** and choose one or more files. Large files (above 8 MB) go directly from your browser to cloud storage, so there is no practical size limit. Each file has **Download** and **Delete**; select several to **Download as ZIP** or delete them together.

> **IMPORTANT**
> Cloud storage has no recycle bin. A deleted file is gone.

Files written on an appliance into `/mnt/datasets/<name>` appear in the Files tab, and files uploaded in the browser appear on the appliance — it is the same storage.

Using a dataset on an appliance

Choose the datasets when you create the appliance (step 4 of the wizard), or later on the appliance's Datasets tab. Each appliance can reach exactly the datasets selected for it and nothing else — an appliance for one project cannot read another project's data. This works on every cloud; outside AWS, access is slower and reading is billed as transfer (see chapter 6).

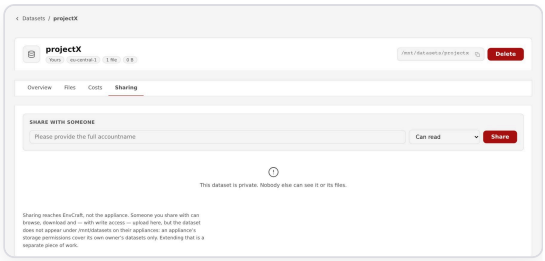


An appliance's Datasets tab, with the notice shown on Google Cloud

Sharing a dataset

The simplest way to share is a project: its members reach its datasets at their level (chapter 5). To share one dataset with one person:

1. Open the dataset's Sharing tab.
2. Enter your colleague's full account name (their email address).
3. Choose **Can read** or **Can read and write**, and click **Share**.



Your colleague can then browse and download the files — and upload, with write access — and select the dataset for their own appliances, where it appears under `/mnt/shared/`. **Revoke** removes the access again; running appliances are reconfigured automatically.

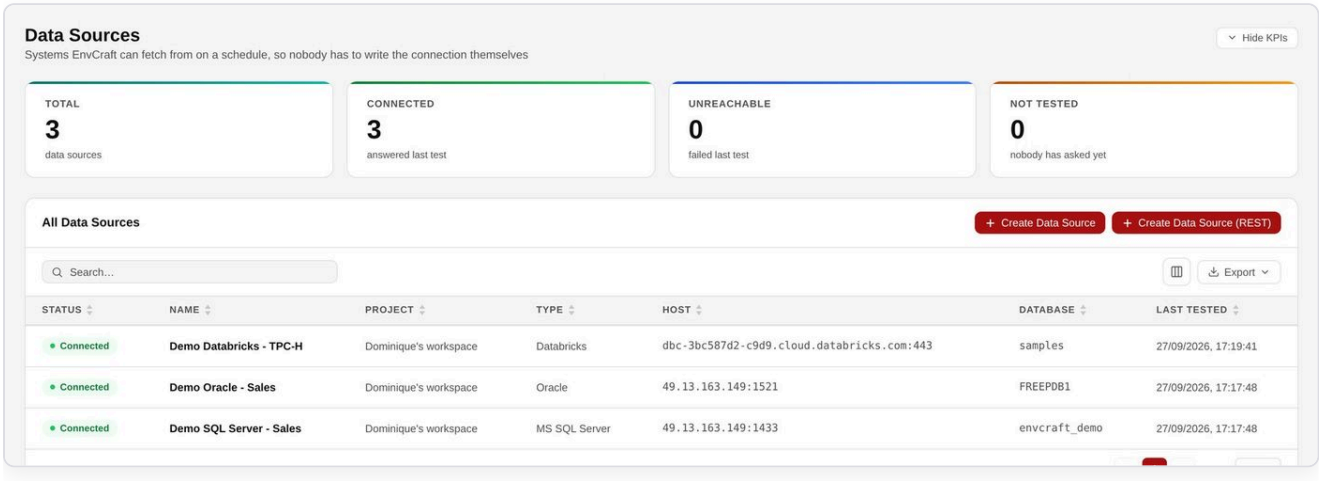
Deleting a dataset

The owner, or a project Admin, can delete a dataset: **Delete** in the workbench header, then **Delete dataset**. All files are removed, all shares revoked, and appliances that mounted it are reconfigured. This cannot be undone.

9 Data sources and scheduled transfers

Connect a system once, write a query once, and let EnvCraft deliver fresh results into a dataset on a timetable.

The flow is always the same: data source → saved query → schedule → dataset → your appliance reads the file from `/mnt/datasets`.



The data source list: Databricks, Oracle and SQL Server connections, each tested before it was saved

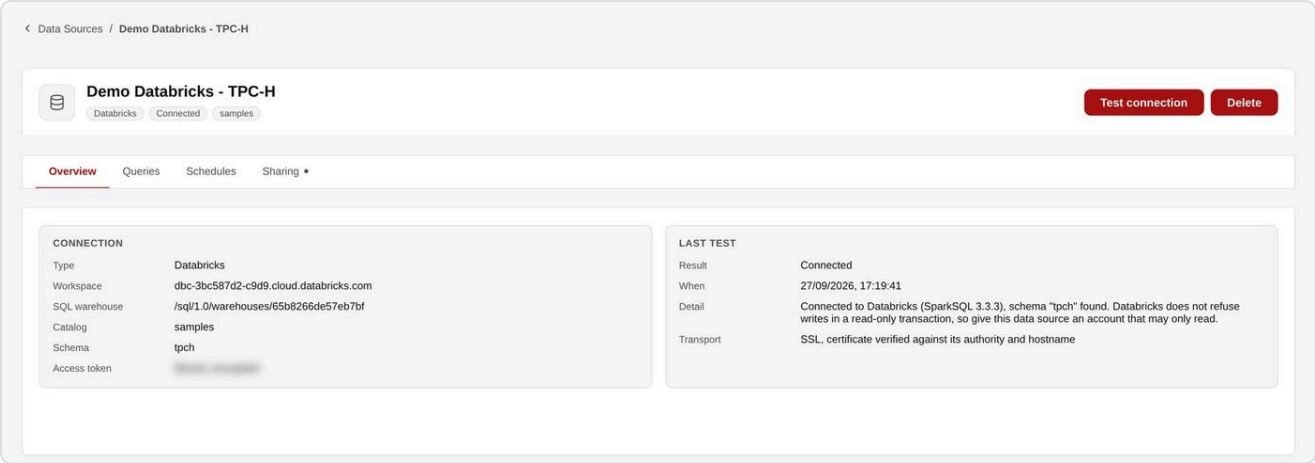
Connecting a database

Click **Create Data Source** and choose the kind. Working today: **PostgreSQL**, **Amazon Redshift**, **Microsoft SQL Server**, **Oracle** and **Databricks**. SAP HANA (VDIL), Denodo, SharePoint and Cassandra are shown as coming soon.

1. Give the connection a **Name** and choose its project.
2. For a database server, enter **Host**, **Port** (the kind's default is filled in), **Database**, **Schema**, **Username** and **Password**, and decide on **Use SSL** and **Verify certificate**.
3. For **Databricks**, enter the **Workspace URL** and an **Access token**, click **Find warehouses** and choose the **SQL warehouse** to run queries on; then the catalogue and schema.
4. Click **Test connection** — EnvCraft really connects and shows the result.
5. Click **Create Data Source**.

Passwords and tokens are encrypted before they are stored and are never shown again.

> **IMPORTANT**
> **Use a database account that may only read.** EnvCraft runs every query read-only and rolls it back. PostgreSQL itself refuses a write in such a transaction; Redshift, SQL Server, Oracle and Databricks do not enforce that in the same way, so only a read-only account really protects the data. The connection test says so.



A Databricks data source: workspace, SQL warehouse, catalogue and schema, and the last test

Connecting an open-data API (REST)

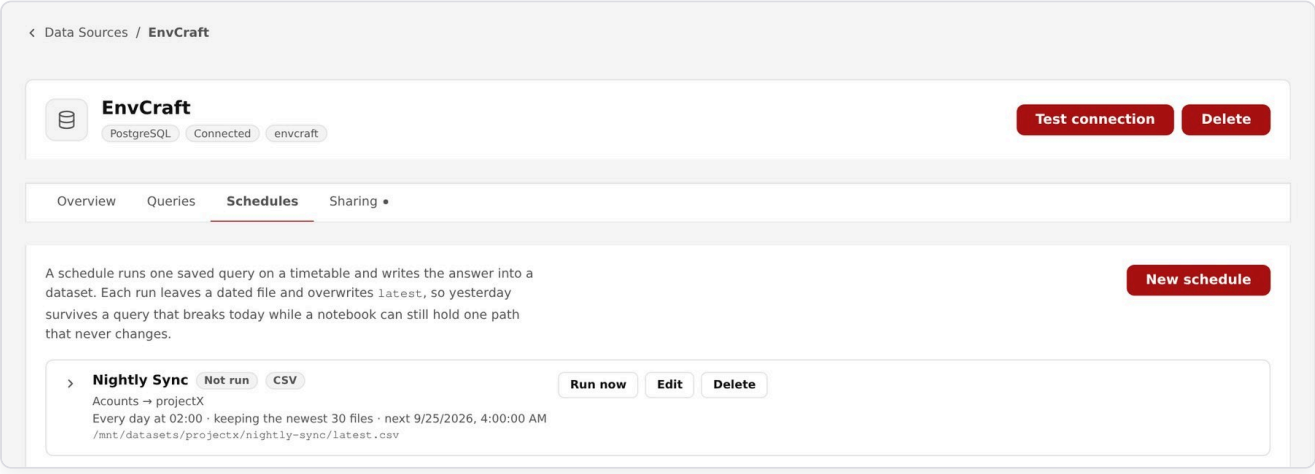
Click **Create Data Source (REST)**, pick an API from the catalogue, give it a name and confirm. Available today: World Bank, ECB Exchange Rates, NYC Open Data, Open-Meteo, NHTSA vPIC, openFDA, USGS Earthquakes, GBIF, Open Library and UK Carbon Intensity — open data, no key or account required. Further connectors are shown as planned.

Saving a query

1. On the data source's Queries tab, click **New query**.
2. Write the SQL (or, for an API, choose what to fetch and its parameters).
3. Click **Run query**. EnvCraft shows the first 100 rows.
4. Once it has run without error, give it a **Name** and click **Save query**.

A query has to run successfully before it can be saved — and again after every change — so a schedule never picks a broken query.

Scheduling a transfer



The Schedules tab, with a schedule's run history folded underneath it

On the Schedules tab, click **New schedule** and fill in:

Field	Meaning
Name	How the schedule appears in the list.
Query	One of the saved queries.
Format	CSV (the only format available today).
Dataset and Folder	Where the files are written — a dataset you may write to, and a folder inside it. The folder is fixed once the schedule exists.
When	Every hour, every day, every week (on a weekday) or every month (on day 1–28, so it also runs in February), at a chosen time.
Keep how many files?	You must choose: **Keep every file**, or **Keep the newest…** N files. Older files are removed only after a successful run.
Run on this timetable	Untick to pause the schedule without deleting it.

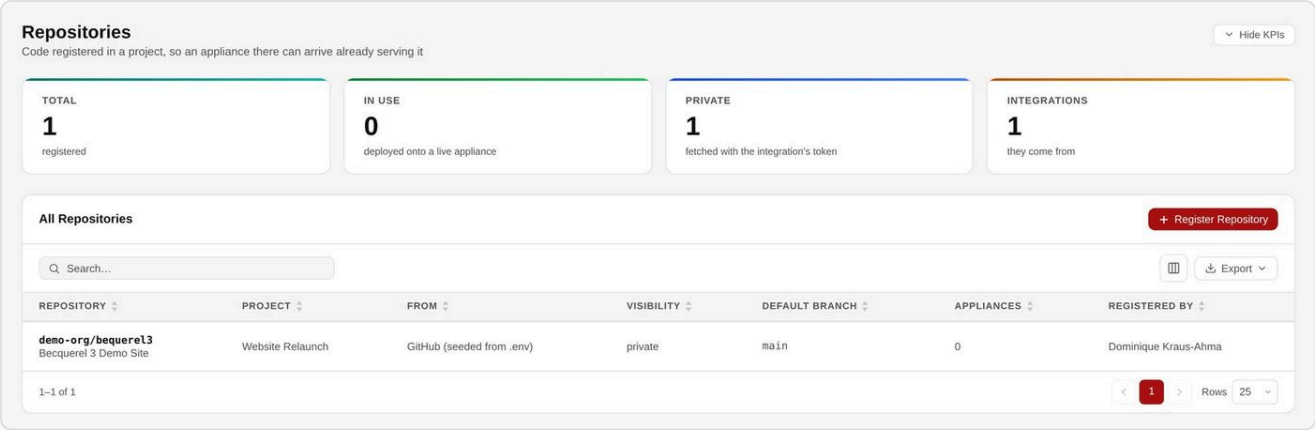
Each run writes a dated file such as `2026-09-24T020000.csv` and overwrites `latest.csv`. Point your notebook at `latest.csv` for the newest data; the dated files let you go back if a run brings a surprise. **Run now** starts a run immediately. Click a schedule to expand its run history.

NOT BUILT YET

The Sharing tab of a data source is not built yet; share it through a project instead.

10 Repositories and deploying code

Register a GitHub repository in a project, and anyone with Write can create a web server that is already serving it.



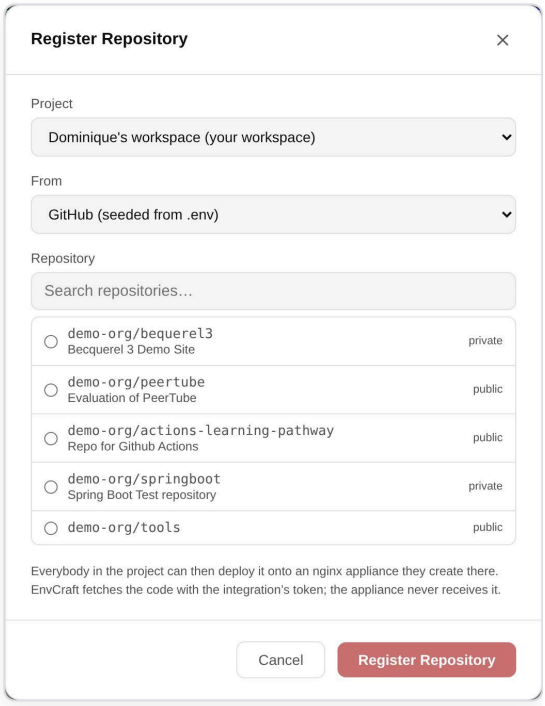
The Repositories page: what is registered in your projects, and how often it is deployed

Registering a repository

Your administrator first connects GitHub as an integration (chapter 14). Then:

1. On the Repositories page, click **Register Repository**.
2. Choose the **Project** and where the repository comes **From** (the GitHub integration).
3. Search for the **Repository** among those the integration's token can see, and select it.
4. Click **Register Repository**.

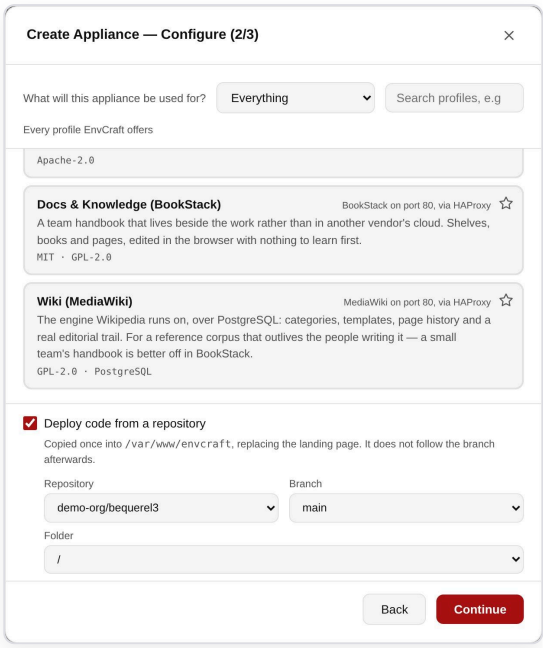
Registering needs Write in the project; removing a repository needs the project's Admin, or the person who registered it. Its workbench has two tabs: **Overview** (what it is and where it is registered) and **Appliances** (which appliances it was deployed onto, and at which commit).



Deploying onto a new web server

1. Create an appliance in the project with the **nginx (web server)** template.
2. In step 3, tick **Deploy code from a repository**.
3. Choose the **Repository**, the **Branch** and the **Folder** that holds the site.
4. Continue and provision as usual.

While the machine is built, EnvCraft fetches that folder at the branch's current commit and copies it into `/var/www/envcraft`, replacing the landing page.



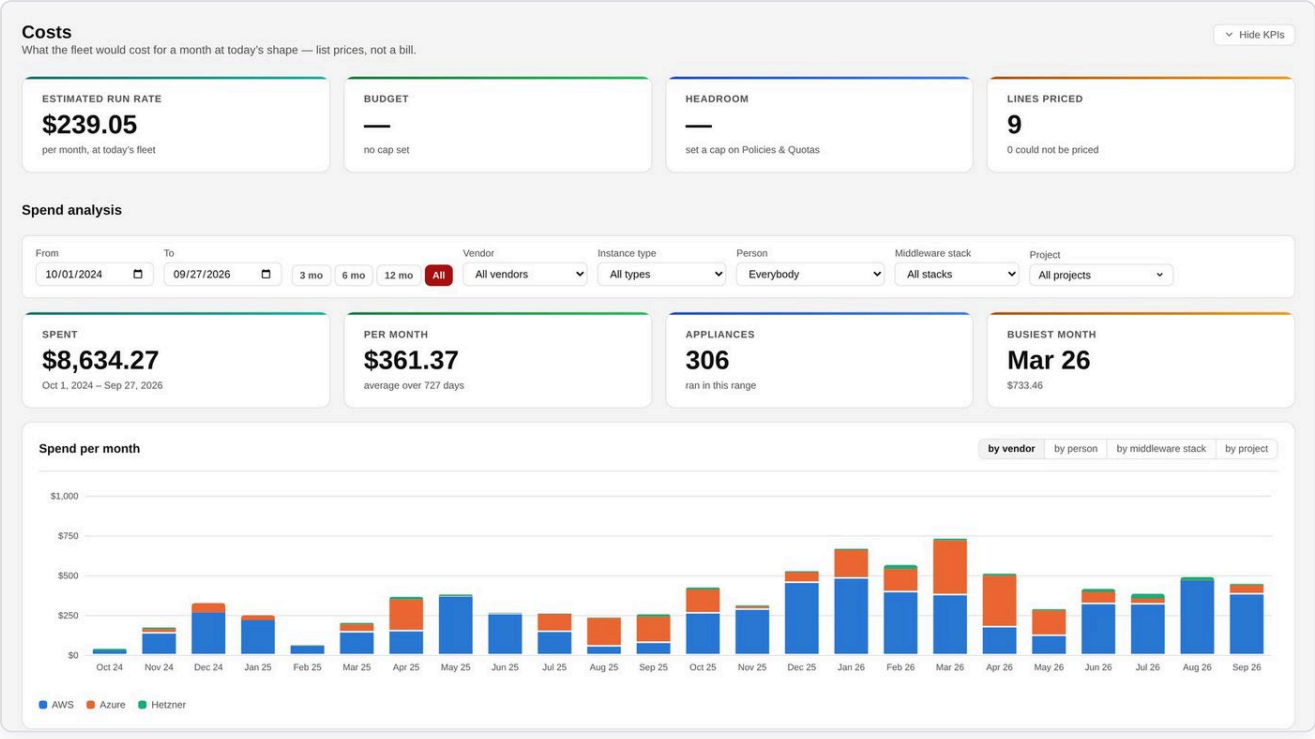
- **The token never reaches the machine.** EnvCraft downloads the code with the company's integration token and hands the machine only the files; the appliance cannot reach the repository afterwards.
- **Deployed once.** The site is pinned to the commit it was built from; nothing follows the branch afterwards. To publish a new version, create a new appliance — or update the files over SSH.
- If the deploy fails, the appliance still starts, with the landing page and the reason on its Logs tab.
- Today this works for the nginx template, with repositories from GitHub.

A build runner for a repository

The **Build Runner** template turns an appliance into a self-hosted GitHub Actions runner for one registered repository. EnvCraft asks GitHub for a one-hour registration token, and only that reaches the machine; jobs run as their own user without administrator rights, and the runner is removed from the repository when the appliance is destroyed. The integration's token needs permission to administer the repository's runners.

11 Costs

What your cloud environments cost — now, per item, and over any period, sliced any way you need.



The top of the Costs page: run rate and budget, then the spend analysis with its filters and key figures

Now: run rate and budget

- **Estimated run rate** — what everything currently running would cost per month.
- **Budget** and **Headroom** — your company's monthly cap, if one is set, and how much room is left.
- **Lines priced** — how many items could be priced, and how many could not.

While the page works out its figures, the tiles shimmer and a short checklist shows which of the answers is still being calculated.

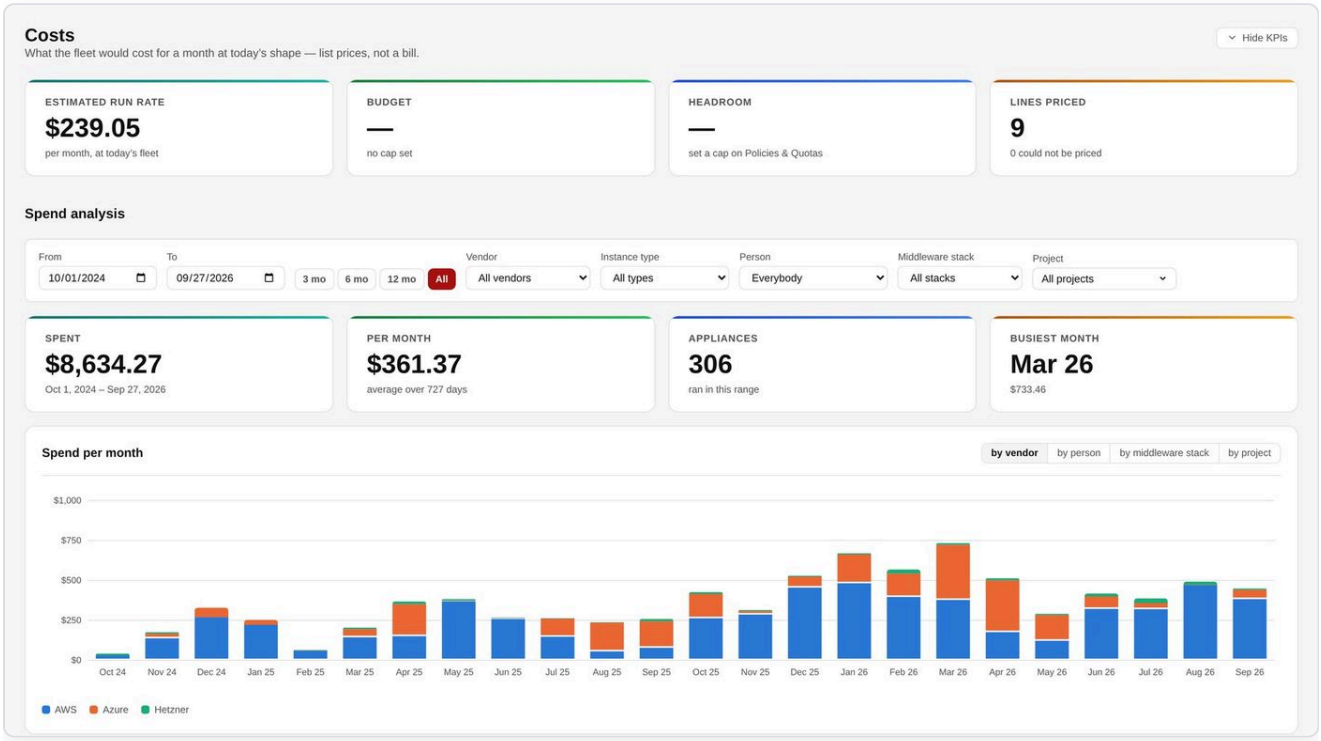
The spend analysis

The larger part of the page analyses what was actually spent, month by month, including appliances that have since been destroyed. One row of filters scopes everything below it:

Filter	Choices
From / To	Any period, with the shortcuts 3 months, 6 months, 12 months and All.
Vendor	All vendors, or AWS, Azure, Google Cloud, Hetzner.
Instance type	All types, or one machine type.
Person	Everybody, or one owner.
Middleware stack	All templates, or one.
Project	All projects, or one or several — searchable.

For the selected slice you see four figures — **Spent**, the average **Per month**, how many **Appliances** ran and the **Busiest month** — and these charts:

- **Spend per month**, stacked **by vendor**, **by person**, **by middleware stack** or **by project**. Every bar has a tooltip; each vendor, person, stack or project keeps its colour.
- **Spend by person** and **Appliances running per month**.
- **By vendor**, **By middleware stack** and **By instance type** — breakdowns with amounts and shares. Click a row to filter by it.
- **Appliances, past and present** — the same slice as a table, with cost and hours inside the range, so its rows add up to the charts.



Monthly spend split by project

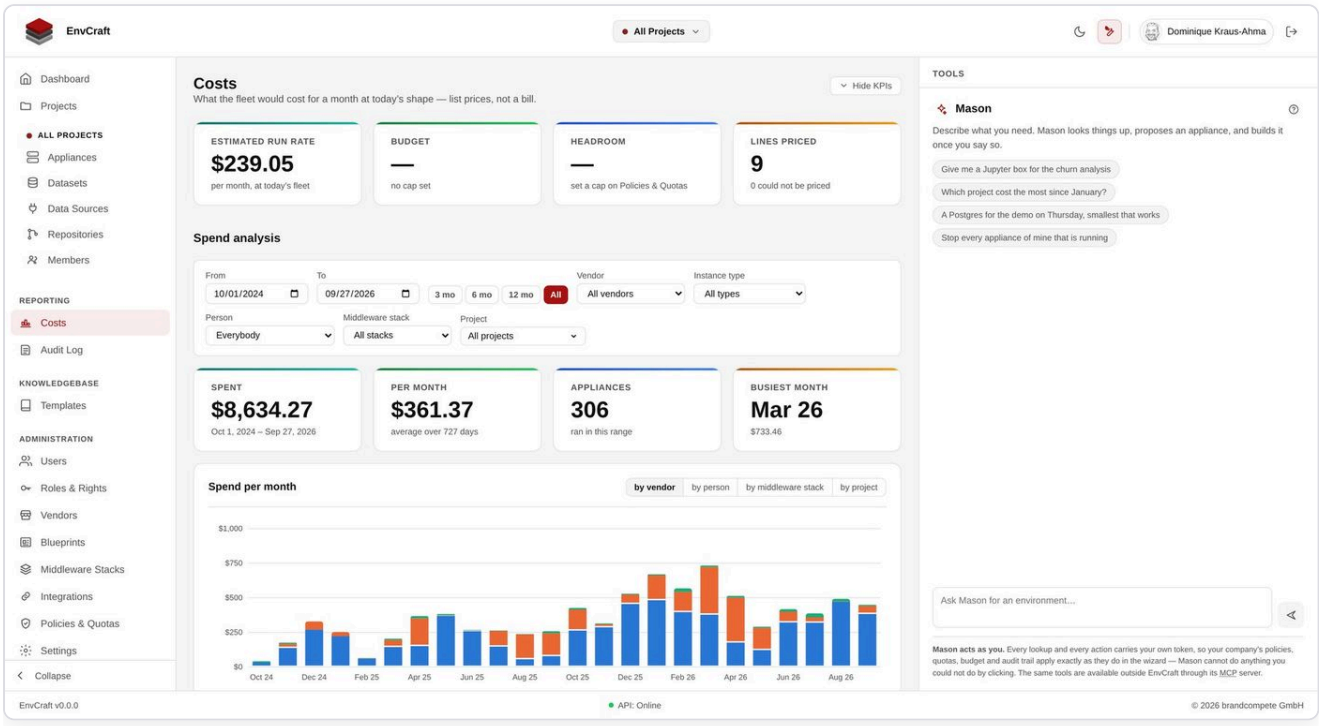
Each project has the same view on its own **Costs** tab, narrowed to what ran in it. **Cost Lines** at the bottom of the page lists today's individual items. Without company-wide rights you see your own items only.

NOTE

All figures are estimates from the providers' public price lists, multiplied by what each item holds. They are not your bill: discounts, reserved capacity, credits, free allowances and data transfer are not included, and dataset storage is not counted per project. Hetzner's euro prices are converted to US dollars at the ECB reference rate; Google Cloud is priced at list price without sustained-use discounts. Items EnvCraft cannot price are marked “not priced” rather than counted as free.

12 Mason, the assistant

Describe what you need in your own words. Mason looks things up, proposes an appliance, and builds it once you say so.



Mason in the Tools panel, opened from the header

Asking Mason

1. Click the **Tools** button in the header, between the theme switch and your avatar. The panel opens on the right, beside the page you are on.
2. Type your request into **Ask Mason for an environment…** and send it — or click one of the suggestions.
3. Mason shows each lookup it makes — the vendors, credentials, regions, sizes, datasets, projects or appliances it checked — and then answers or proposes.
4. When Mason proposes something that needs your consent, a button appears, such as **Create this appliance**. Nothing happens until you click it.

What Mason can do

Mason works with the same catalogue of tools as the MCP server — some fifty of them: look up vendors, regions, sizes, prices, credentials, SSH keys and templates; list, create, rename, start, stop and destroy appliances and set their power schedules; create, mount and share datasets; list and create projects and manage their members; list repositories and deploy one onto a new nginx appliance; report costs by project, person, vendor, size, template or month; read the budget, policies and audit log; and change settings and your own preferences, such as the theme.

What always waits for you

Anything that spends money, grants access, reconfigures a machine or changes a setting waits for your confirmation — creating, starting or stopping an appliance, a power schedule, mounting or sharing datasets,

adding or removing members, creating a user or changing a role, company settings and switching a vendor on or off. Destroying an appliance or deleting a dataset over MCP also needs its exact name.

What Mason never gets

Anything that takes a secret (cloud credentials, data source passwords, integration tokens), the terminal, deleting people and editing a role's rights.

Mason acts as you

Every lookup and every action carries your own sign-in, so your company's vendors, instance types, quotas, budget, project levels and audit trail apply exactly as they do in the wizard. Mason cannot do anything you could not do by clicking; a tool you have no right for is not even offered to it, and Mason can tell you which right is missing. An appliance it creates is yours in the audit log.

NOTE

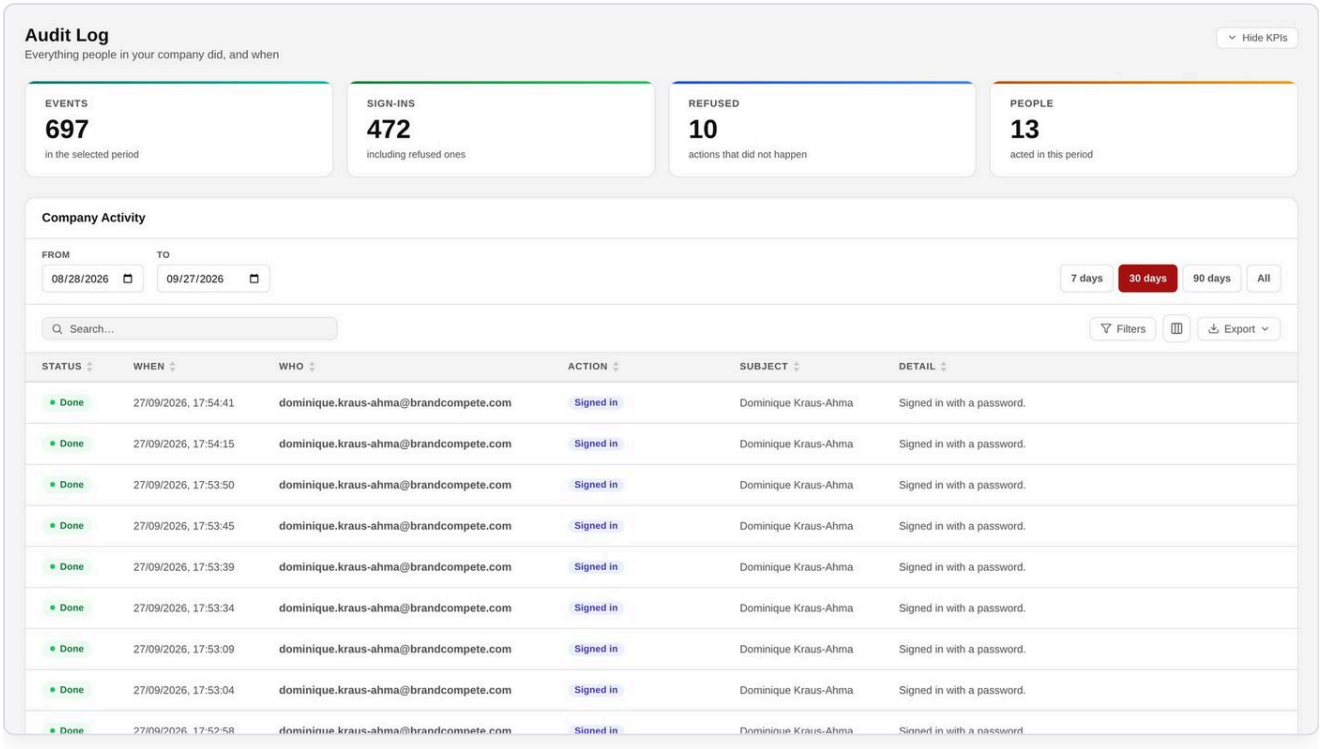
Using Mason needs the right `can_use_assistant`; without it the Tools button does not appear, and Mason answers that your role does not include it. Every message is a paid call to a language model (Anthropic's Claude), made from the EnvCraft server — the key never reaches your browser. An installation that must not call an external model leaves Mason switched off; the panel then shows a **Concept** badge, answers from a fixed script and changes nothing.

The same tools outside EnvCraft

EnvCraft also offers its tools over the **Model Context Protocol (MCP)**, so an assistant such as Claude Desktop, Claude Code or your own MCP client can work with EnvCraft. The MCP server signs in as one EnvCraft account and acts with that account's rights; your administrator should give it an account and role of its own, so that the audit log shows who did what. Ask your administrator if you want to use it.

13 The audit log

Who did what, and when.



The audit log

The **Audit Log** records sign-ins (including refused ones), appliances created, powered and destroyed, terminal sessions opened and closed, datasets created and shared, projects and their members, credentials and SSH keys added or removed, integrations and role changes — and every action EnvCraft refused.

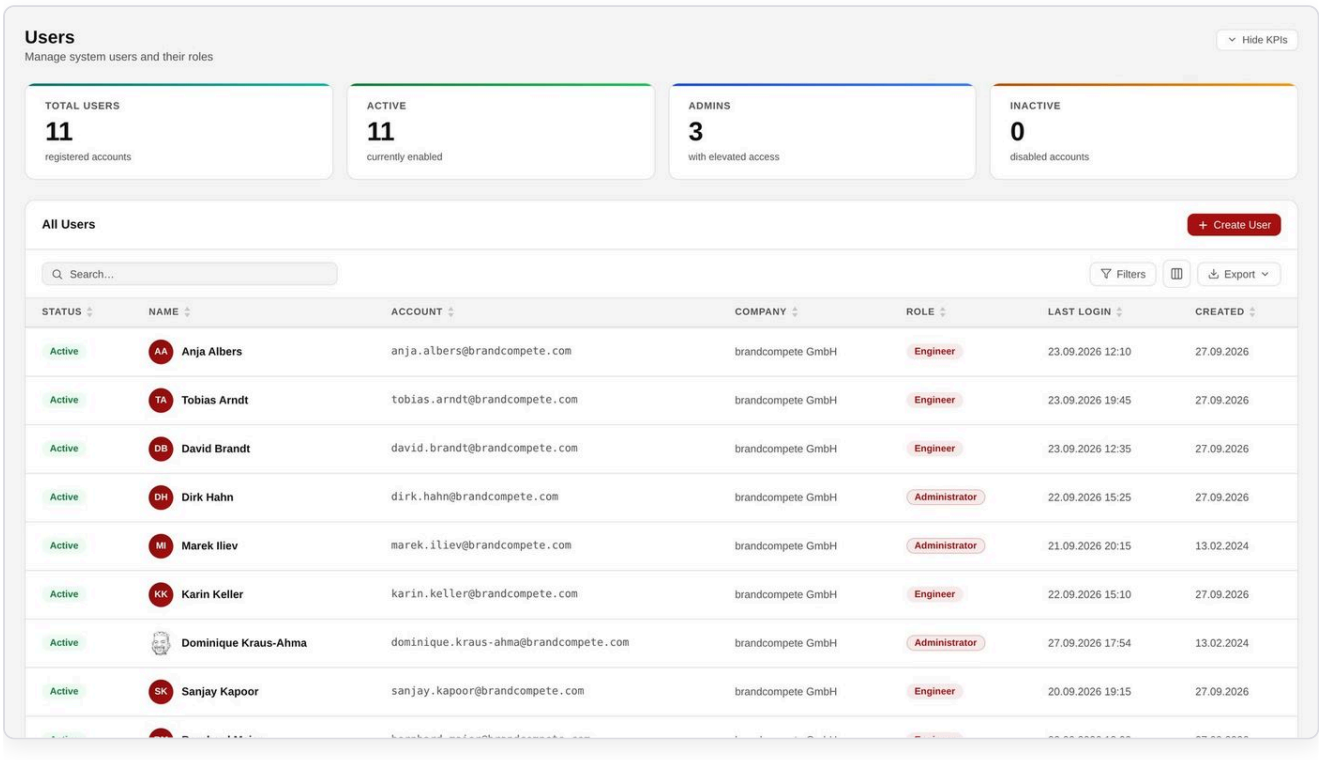
- Everyone sees their own events. People with the right to view the audit log see the whole company.
- Choose a period with **From** and **To** or the buttons **7 days**, **30 days**, **90 days**, **All**.
- Filters narrow by kind of event and by outcome (**Done** or **Refused**).
- Export downloads the selection as Excel, CSV or PDF.

Administrators find the same trail, narrowed to one person, on that person's Activity tab. Actions taken through Mason or the MCP server appear under the person they were taken for.

14 Administration

For administrators: people, roles, clouds, integrations, the template catalogue, limits and company-wide defaults. You see these pages only if your role carries the rights for them.

Users



The user list

The Users page lists the people in your company with their status, role and last sign-in.

Creating a user

1. Click **Create User**.
2. Enter first name, last name and email address. EnvCraft suggests an initial password (at least 8 characters) — you can change it.
3. Choose a **Role** (default: Guest) and whether the person may use the **Web interface**, the **API**, or both.
4. Click **Create User**, then give the person their password yourself — EnvCraft sends no email.

New users always join your own company and get a personal workspace. People who sign in with Microsoft or Google for the first time are created automatically with the role Guest.

Editing and deleting

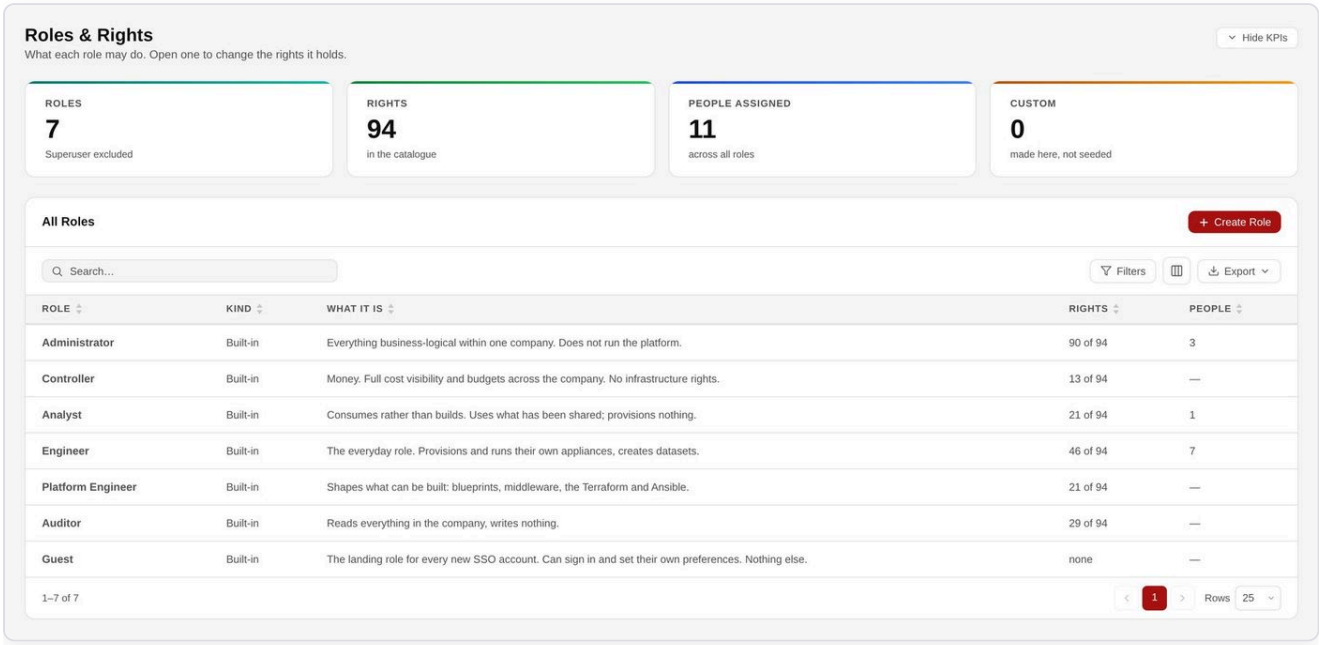
Click a user to open their workbench (Overview • Edit • Activity). On Edit you change their photo, name, role, web/API access and — to help someone who forgot it — set a **New password**. Leave it empty to keep the current one.

Delete in the header deactivates the account; the record stays so the audit log keeps showing who did what. Deleting is refused for your own account, and while the person still owns appliances, datasets, data sources or shares — EnvCraft names what is in the way.

TIP

To offboard someone temporarily, give them the role **Nobody**: they can no longer do anything, but remain easy to find.

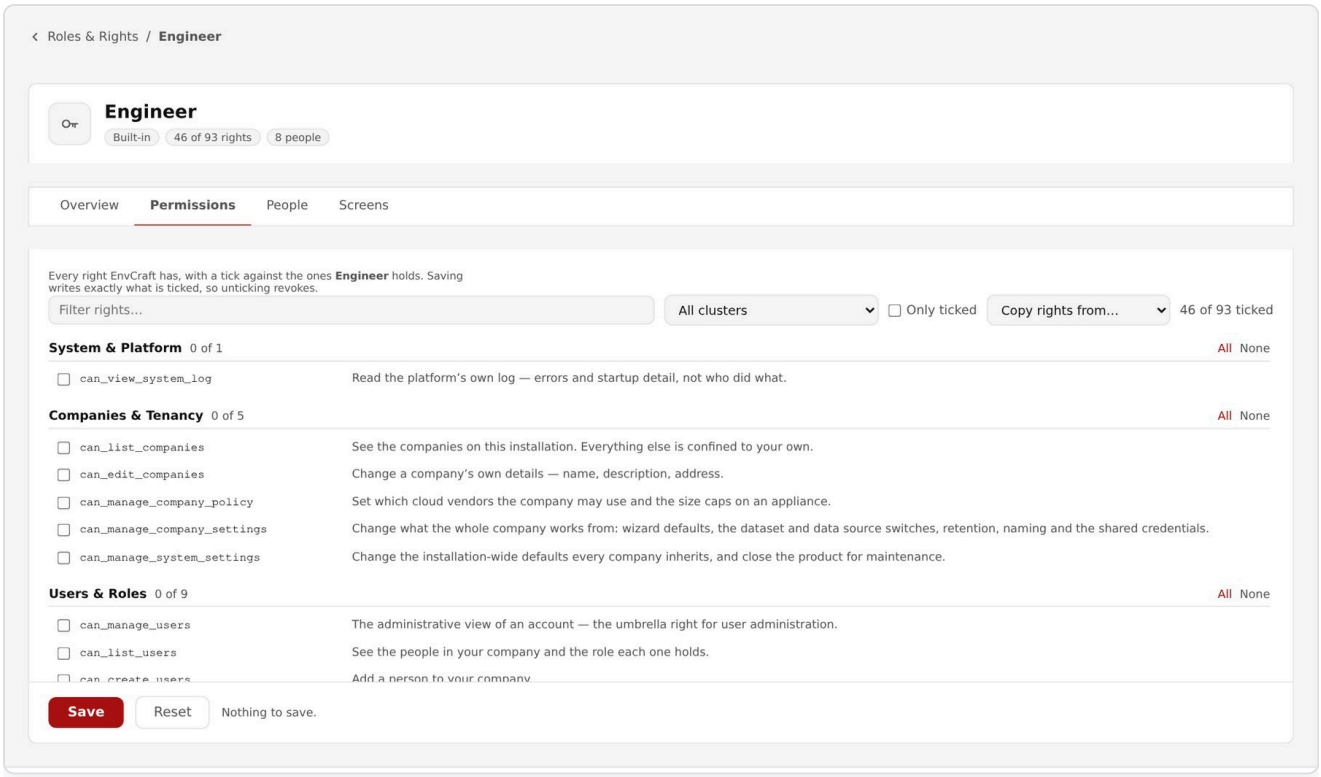
Roles & Rights



The role list: the built-in roles with what each is for, its rights and who holds it

Roles are sets of rights. EnvCraft comes with built-in roles (see Appendix A), and you can create your own — for example “Contractor”, who may read datasets and nothing else.

1. Click **Create Role**, enter a **Name** and what it is for. The new role starts without any rights.
2. Open it and go to **Permissions**. Tick the rights it should carry, grouped by area; each right has a description. **Copy rights from…** loads another role's rights as a starting point.
3. Click **Save**. A summary shows exactly which rights are granted and revoked, and how many people are affected, before anything changes.



The Permissions tab

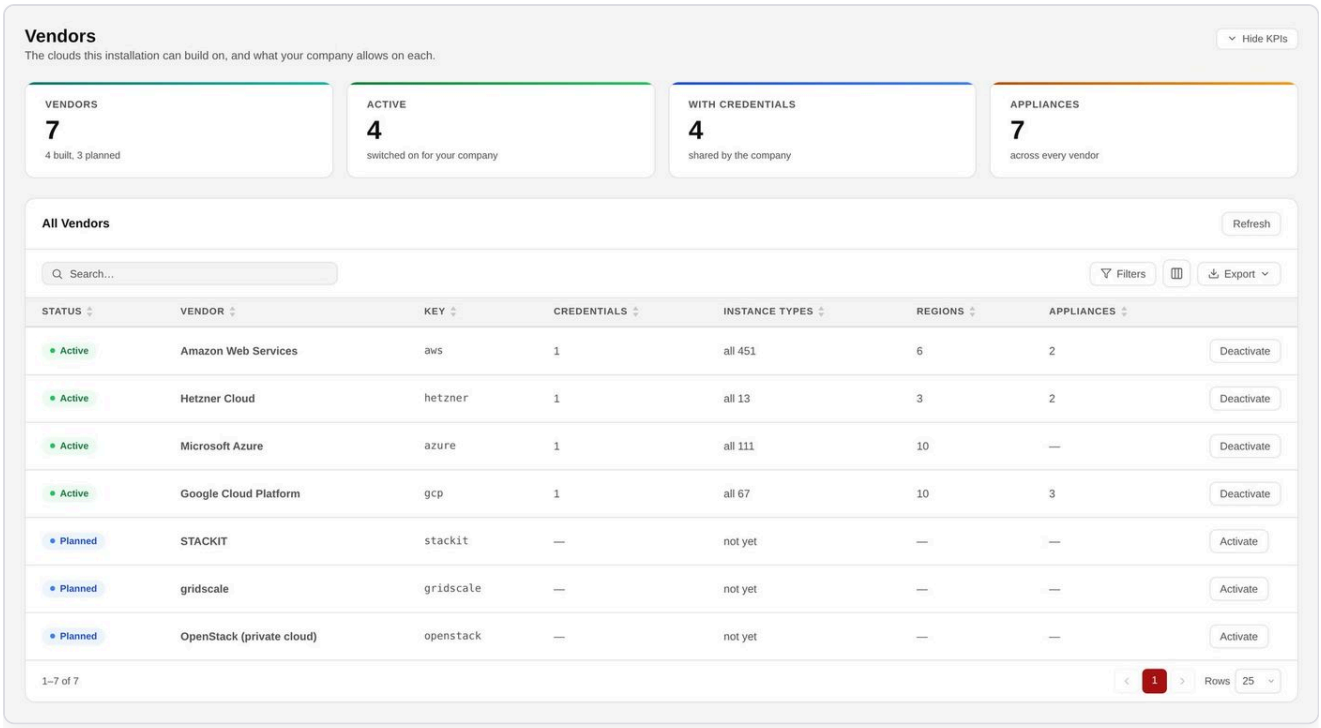
The **People** tab lists who holds a role; **Screens** shows which parts of the application the role opens. Built-in roles cannot be renamed or deleted, but their rights can be adjusted. A custom role can be deleted only when nobody holds it any more. Remember that project levels come on top: a role opens a kind of action, a project decides where.

IMPORTANT

Roles apply to the whole installation. If EnvCraft serves several companies, changing a role changes it for everyone who holds it — the People tab shows which companies they belong to.

Vendors

The clouds this installation can build on, and what your company allows on each. There is no Create button: a vendor exists because EnvCraft can build on it — or plans to.



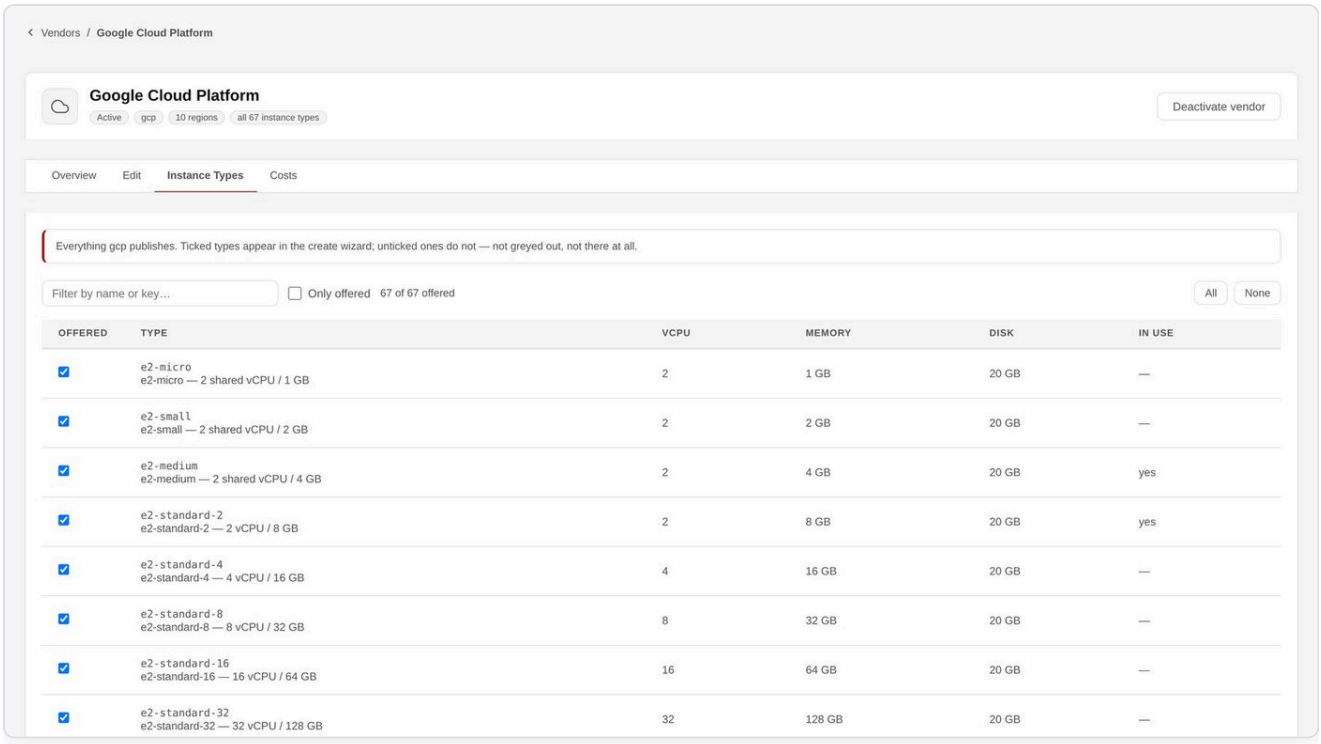
The vendor list: four vendors built, three planned

The list shows each vendor's status, its credentials, how many instance types and regions it offers, and how many appliances run on it. Amazon Web Services, Hetzner Cloud, Microsoft Azure and Google Cloud Platform are built; STACKIT, gridscale and OpenStack (private cloud) are **Planned**.

Status	Meaning
Active	Built, and switched on for your company — offered in the create wizard.
Inactive	Built, but switched off for your company. Its card in the wizard is greyed out.
Planned	EnvCraft cannot build on it yet. Off by default.
Active • no adapter	A planned vendor switched on anyway, for example to show a prospect their own cloud in a demonstration. The wizard says that provisioning will not succeed.

Activate / **Deactivate** in the list or the workbench header switches a vendor on or off for your company. Click a vendor to open its workbench:

Tab	What you find there
Overview	Availability and what your company uses on this vendor.
Edit	Whether appliances may be created on this vendor, and the **company's cloud credentials** for it — offered to everyone in the wizard as "(company)". More than one per vendor is normal. **+ Add credential** uses the same form as personal credentials (chapter 4).
Instance Types	Everything the vendor sells in the region, read live. Ticked types appear in the create wizard; unticked ones do not — not greyed out, not there at all. **Only offered**, **All** and **None** help with long lists; **In use** shows types that appliances run on. A type the vendor launches tomorrow is offered automatically.
Costs	This vendor's share of the company's spend.



Google Cloud's Instance Types tab

What works where	AWS	Azure	Google Cloud	Hetzner
Appliances, live instance types, prices	yes	yes	yes	yes (EUR → USD)
Power schedules	yes	yes (deallocates)	yes	no
Datasets	yes	yes, own key	yes, own key	yes, own key
Monitoring	yes	not yet	not yet	not yet

NOTE

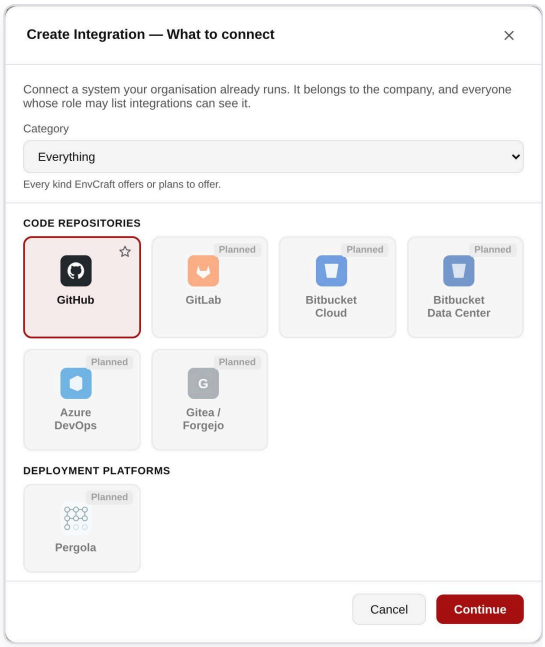
Neither switch reaches back: an appliance already running on a vendor you deactivate, or on a type you withhold, keeps running and keeps being priced. To stop it, destroy it from its own workbench.

Integrations

Connections to the tooling your organisation already runs, configured once for the whole company. **GitHub** is available today; GitLab, Bitbucket, Azure DevOps, Gitea/Forgejo, the Pergola deployment platform, container, package and model registries, secret stores and notification channels are listed as planned.

Connecting GitHub

1. On the Integrations page, click **Create Integration**.
2. Pick **GitHub** from the catalogue. The **Category** drop-down narrows it; the star keeps a kind among your favourites.
3. Enter a **Name**, the **Personal access token** and optionally a **Description**.
4. Click **Test connection** — EnvCraft asks GitHub who the token belongs to — then **Create**. The token is checked with GitHub again before anything is stored.



The catalogue: GitHub available, the rest planned

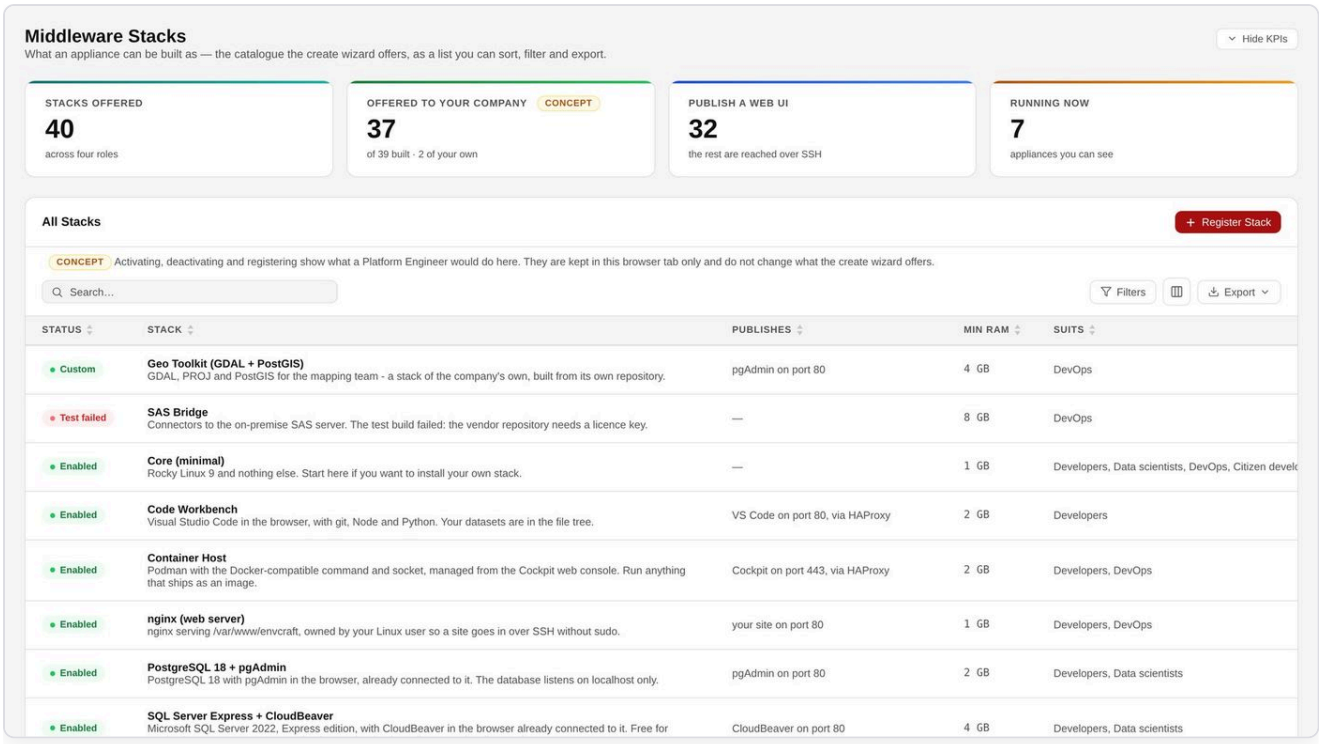
IMPORTANT

A company integration acts on GitHub as whoever the token belongs to, for everybody who uses it. Prefer a fine-grained token on a machine account, limited to the repositories that matter, with read access to *Contents* and *Metadata* — plus *Administration* on repositories that get a build runner. The token is encrypted and never shown again, and it never reaches an appliance.

The list shows each integration's status — **Connected**, **Failing**, **Not tested** or **Token expired** — its type, the account it acts as, when the token expires (amber within three weeks, red once it has lapsed) and when it was last tested. Its workbench has four tabs: **Overview**, **Repositories** (what the token can see, asked of GitHub live), **Edit** (name, description, replace the token) and **Usage** (the deployments made with it).

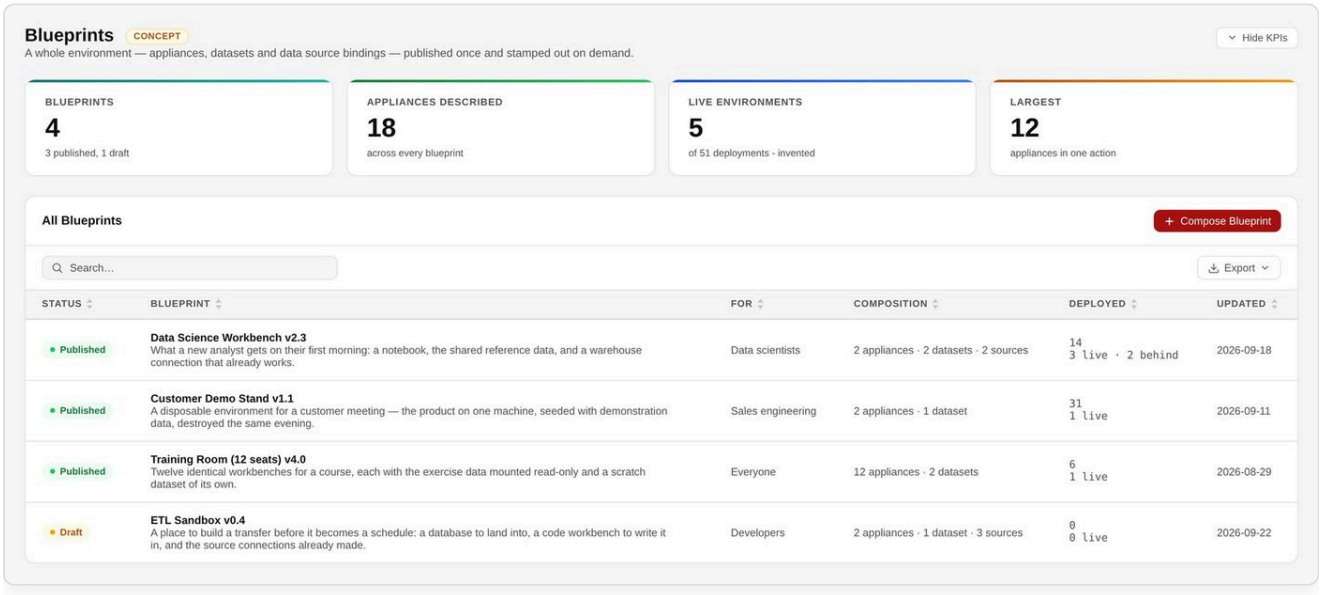
Middleware Stacks and Blueprints

Two screens show how the template catalogue and whole environments could be managed by a **Platform Engineer**. Parts of them are real, parts are a concept.



Middleware Stacks: the real catalogue, with a Concept strip for what a Platform Engineer would do here

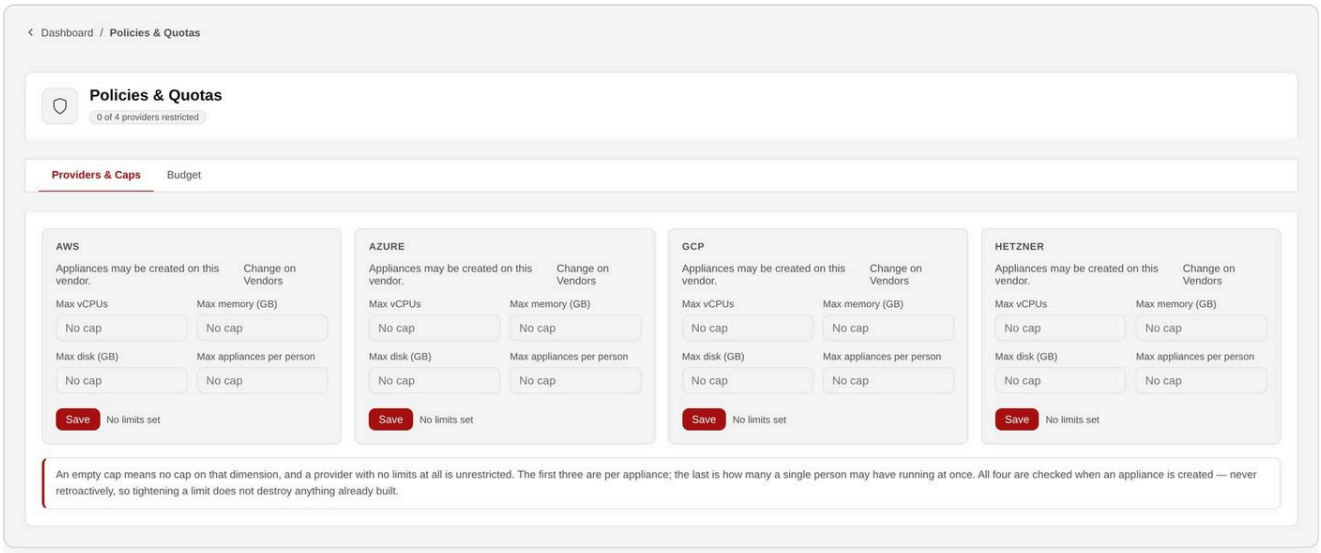
- **Middleware Stacks** — real: the catalogue, how many appliances run each template, and each template's Overview, Appliances and Provisioning tabs. **Concept** : enabling and disabling templates for the company, **Register Stack** for a company's own template from a repository, and the Configuration and Versions tabs. These are kept in your browser tab only and do not change what the create wizard offers.
- **Blueprints** **Concept** — a whole environment of appliances, datasets and data source connections, composed once and deployed in one action. Composing works with the real catalogue, and **Deploy** runs a plan — the real checks in the real order — without creating anything. Nothing survives a page reload.



Blueprints, a concept

Policies & Quotas

Limits that EnvCraft checks before anything is created in a cloud. A request that would break one is refused with a message naming the limit and the current value. Tightening a limit never affects appliances that already exist.



Providers & Caps, per cloud

Providers & Caps

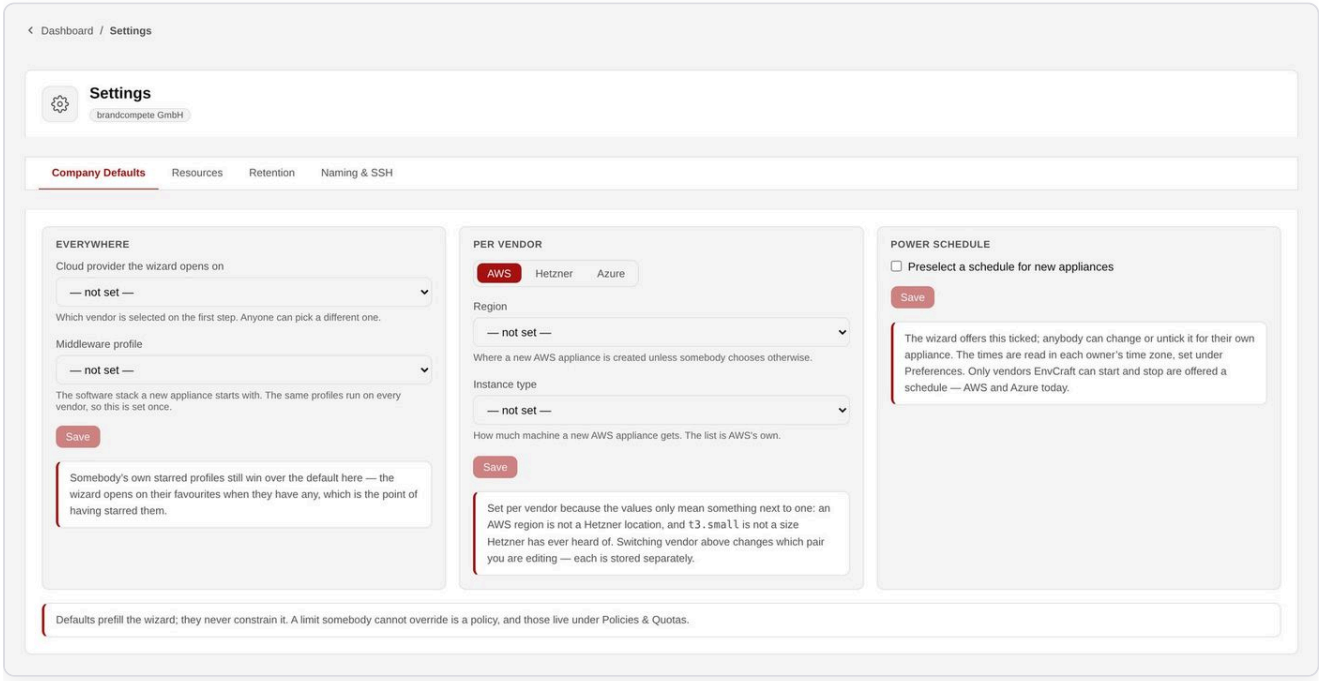
Per cloud vendor: the maximum vCPUs, memory and disk per appliance, plus appliances per person. An empty field means no cap. Whether appliances may be created on a vendor at all is switched on the Vendors page, and each panel links there.

Budget

A monthly **Cap**, measured across the whole company or per person, and what happens above it: **Warn only** or **Refuse new appliances**. The budget is compared with the estimated run rate described in Costs. **Withdraw** removes it.

Company settings

Settings holds your company's defaults. Each panel has its own **Save** button; a field left empty inherits the installation's default and says so.



Company Defaults, with the preselected power schedule

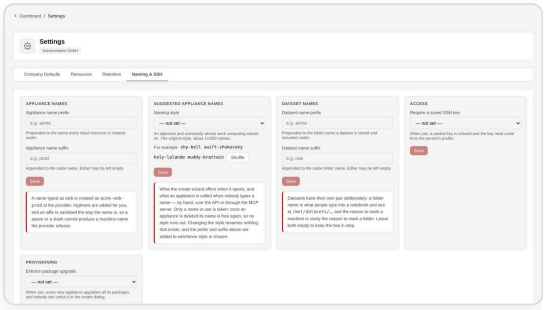
Tab	What you set
Company Defaults	**Everywhere**: the cloud vendor the wizard opens on and the default template. **Per vendor**: the default region and instance type. **Power schedule**: a schedule the wizard preselects for new appliances. Defaults pre-fill the wizard — they never restrict it. Limits belong in Policies & Quotas.
Resources	Whether datasets are enabled, whether they may be shared, how many each person may have; whether data sources are enabled.
Retention	How long deleted records and audit events are kept. **Stored, not yet applied**
Naming & SSH	A prefix and suffix added to every appliance's and every dataset's technical name (for example your cost centre); the **style of suggested appliance names**; whether a saved SSH key is required — a pasted key is then refused; and under **Provisioning**, whether every new appliance must upgrade all packages. When that is enforced, the box in the wizard is fixed.

Suggested appliance names

When nobody types a name — in the wizard, through Mason or over the API — EnvCraft suggests one in the style you choose:

- **Pioneers** — `brave-hopper` (the default)
- **Cinema** — `jolly-ripley`, curated film characters: no villains, no real people
- **Snack bar** — `grumpy-pretzel`
- **Bureau** — `interim-bureau-of-lost-socks`

The panel previews suggestions in the chosen style; **Shuffle** shows others. Only a live appliance holds a name, so no style runs out.



15 Troubleshooting & FAQ

My appliance shows *Error*.

Click **View log** in the list (or open the Logs tab). The end of the log says what failed. If you cannot resolve it, click **Copy log** and send it to your administrator or brandcompete support. Destroy the failed appliance before trying again, so it does not keep costing money.

My appliance shows *Not in cloud*.

The cloud provider no longer knows the machine — it was deleted outside EnvCraft. Open it and click **Remove** to clear the entry.

The “Continue” button in the wizard stays grey.

Step 2 needs a project, a cloud credential, a region, a size and an SSH key, and a power schedule — if you ticked one — needs at least one weekday and a time. If no credential is offered, add one in your preferences or ask your administrator for a company credential.

The size I want is not in the list.

The list shows what the vendor sells in the chosen region, less the types your company has withheld on the Vendors page. Try another region, or ask your administrator.

I cannot find an appliance, dataset or data source I created.

Look at the project switcher in the header: the lists show only the chosen project. Choose **All Projects**.

“Deploy code from a repository” does not appear.

It is offered only for the nginx template, and only when a repository is registered in the appliance's project (chapter 10).

Provisioning was refused with a message about a limit.

Your company's policies do not allow the requested vendor, size, number of appliances or budget. Choose a smaller size, destroy an appliance you no longer need, or ask your administrator.

My appliance was stopped although I did not stop it.

Look at its Schedule tab: a power schedule stops it at the set time, in the owner's time zone. Start it again by hand if you need it — it stays on until the next stop time.

The terminal icon is grey.

Hover over it for the reason: your role lacks the right to connect, you need Write in the appliance's project, the appliance is not running, or it was built before the terminal existed and must be rebuilt.

I cannot log in with my own SSH client.

- Use the command from **Copy ssh command** — it contains the right user and address.
- Your SSH client must use the private key belonging to the public key chosen when the appliance was created. Keys added later are not on existing machines.
- After several failed attempts, fail2ban blocks your address for a while. Wait, or use the terminal in the browser.

A file I uploaded does not show up on the appliance.

Wait about 30 seconds — listings are cached briefly. Check that the dataset is selected on the appliance's Datasets tab.

My browser asks for a user name and password when I open a notebook or app.

That is the appliance's front door. Use the Linux user and the password shown under Overview › Access.

My browser warns about the certificate when I open Cockpit.

Cockpit on the Container Host uses a self-signed certificate. Accept the warning once for that appliance.

A menu entry or a button is missing or grey — or there is no Tools button.

Your role, or your level in the project, does not allow it. Grey buttons explain what is missing when you hover over them. Ask your administrator or the project's Admin.

I forgot my password.

Ask your administrator to set a new one on your user page. If you normally sign in with Microsoft or Google, use that button instead.

Why is the cost different from our invoice?

EnvCraft shows estimates from public list prices. Your invoice reflects discounts, credits, taxes and data transfer that EnvCraft does not know about — and Hetzner invoices in euros, while EnvCraft converts to dollars at the ECB rate.

A Reference

Built-in roles

Role	Meant for
Administrator	Everything business-related within one company: people, roles, vendors, integrations, policies, settings, and all resources.
Engineer	The everyday role: provisions and runs appliances, creates datasets and data sources, opens terminals.
Platform Engineer	Shapes what can be built: blueprints, middleware stacks and the scripts behind them.
Analyst	Works with what is shared with them; provisions nothing.
Controller	Money: full cost visibility and budgets across the company, no infrastructure rights.
Auditor	Reads everything in the company, changes nothing.
Guest	The starting role for every new single-sign-on account: can sign in, read the Knowledgebase and set their own preferences, nothing else.
Nobody	For offboarding: the account exists but may do nothing.

Together they draw on a catalogue of over ninety rights; whether a role may use Mason is one of them. Your administrator may have added custom roles, and may have changed which rights the built-in roles hold.

Appliance statuses

Status	Meaning
Provisioning	Being created and configured. Watch the Logs tab.
Running	Ready to use.
Stopping / Starting	A power change is under way — by hand or by a power schedule.
Stopped	Powered off. Disk, address and datasets are kept. On AWS a stopped machine costs only its storage and address; on Azure and Google Cloud its disk and static address; Hetzner bills a stopped server in full — destroy it if you no longer need it.
Deleting	Being destroyed.
Error	Something failed — see **View log**.
Not in cloud	The provider no longer has the machine.

Glossary

Term	Meaning
Appliance	A cloud machine plus the software configured on it.

Term	Meaning
Blueprint	A concept: a whole environment — several appliances, datasets and data-source connections — published once and deployed in one action.
Cloud credential	The access key of the cloud account appliances are created in; personal or company-wide, always stored encrypted.
Data source	A connection to an existing system, used by saved queries.
Dataset	A folder in managed cloud storage, mounted into appliances at `/mnt/datasets` (own) or `/mnt/shared` (shared with you or through a project).
Front door	The small HAProxy on each appliance that serves every web interface on the ordinary web port, and asks for a sign-in where the application has none of its own.
Instance type	A machine size a vendor sells, such as `t3.medium`, `e2-medium` or `cpx12`.
Integration	A company-wide connection to the organisation's own tooling, such as GitHub.
Knowledgebase	The pages that describe every template: its software and licences.
Linux identity	A login name you create appliances with; up to five per person.
Mason	The assistant in the Tools panel. It acts with your own rights.
MCP	Model Context Protocol — the way external AI assistants use EnvCraft's tools.
Personal workspace	Your own project, never shared.
Power schedule	Weekdays and a start and stop time at which EnvCraft powers an appliance on and off by itself.
Project	The workspace resources belong to, with members at Read, Write or Admin.
Provisioning	Creating the machine in the cloud and installing its software.
Repository	A GitHub repository registered in a project, ready to be deployed onto a new web server.
Right / Role	A right allows one kind of action; a role bundles rights and is assigned to people.
Run rate	The estimated monthly cost of everything currently running.
Schedule	A saved query bound to a timetable and a dataset folder.
SSH key	The public half of a key pair that lets you log in to your appliances with your own SSH client.
Template	The software recipe an appliance is built with, also called a middleware stack.
Vendor	A cloud EnvCraft can build on — AWS, Azure, Google Cloud, Hetzner — with the company's credentials and allowed instance types.
Workbench	The detail page of an item, with its actions in the header and tabs below.

Questions or feedback: office@brandcompete.com • brandcompete GmbH • www.brandcompete.com