

A RESEARCH OPERATING SYSTEM

LabsOfScience

One governed record for the whole of a lab's work, from the first hypothesis to the result someone else can reproduce.

ELECTRONIC LAB NOTEBOOK

VERSIONED PROTOCOLS

IMMUTABLE DATASETS

REPRODUCIBLE PIPELINES



Pre-launch. The research system of record and the web application are live; the waitlist is open ahead of the 2026 launch.

LABSOFSOFTWARE.COM

01 The short version

One place for everything a lab actually does: designing an experiment, running it, writing it up, versioning the data, analysing it, and keeping track of the reagents and instruments behind all of it.

Most research groups run on a patchwork: a notebook tool here, a drive of unversioned spreadsheets there, a method emailed round as a PDF, a freezer list in somebody's head. Together they produce the quiet, expensive problem every lab recognises.

LabsOfScience puts the whole lifecycle on a single connected record — and none of it is a discipline the team has to remember.

THE PROMISE

If you can see a result, you can walk back to the exact data, the exact method version and the exact run that produced it, without anyone having curated that chain by hand.

THE PRODUCT IN FOUR FIGURES

One

Connected record

Projects, experiments, runs, notebook entries, protocols, datasets, pipelines, inventory and instruments share a single model instead of syncing between tools.

45

Governed agent tools

The whole research domain is reachable by an AI agent, with six irreversible actions refused unless intent is explicitly confirmed.

11

Permission-checked record types

Every read and write is authorised at the API edge against the caller's role before it reaches the product.

Day 1

Governance included

Single sign-on, role-based access, audit, file storage, tagging, billing and notifications arrive with the product, not as add-on modules.

02 Who it is for

Academic research groups

A real electronic notebook, versioned methods and reproducible analysis, at a price a grant line can carry.

Biotech and life sciences

Provenance, a notebook a second person can witness, and reagent and instrument tracking an auditor will accept.

Materials, chemistry, engineering R&D

Structured experiments with declared variables and physical stock behind them. The model does not assume molecules.

Computational research teams

Groups that version datasets and re-run pipelines, and need the chain from data to result to survive a reviewer.

Core facilities and shared labs

Instrument scheduling and shared consumables: bookings that cannot overlap, stock levels that update themselves.

Research-driven enterprises

Corporate R&D consolidating notebooks, methods, data and inventory into one governed platform with isolation between programmes.

And research and automation engineers — the people who script the lab. Everything in the application is available over the same programmable interface, under exactly the same permissions as the user interface.



03 What breaks today

The ordinary friction of doing science with tools that were never designed to hold a research record together. It compounds quietly until the day somebody asks a question the record cannot answer.

WHERE THE RECORD FAILS

The record is spread across five tools

Design in a spreadsheet, write-up in a notebook product, data on a drive, method as a PDF. Nothing links to anything.

One data model holds all of it — method, data and reagents attached to the same record rather than referenced by memory.

Data files multiply until nobody knows which is final

The file named final, the one named final version two, and the one somebody renamed on a Friday.

A dataset has numbered versions, each with a checksum and a change note. The version a result used is recorded, not remembered.

The method lives in one person's head

The postdoc who knew the assay leaves. New starters reinvent the method and get different answers.

Protocols carry ordered, typed steps and safety notes. Publishing freezes the version, so an experiment names the exact edition it followed.

The analysis only ever ran once, on one machine

When the reviewer asks for the analysis, what gets sent is a folder and an apology.

A pipeline run is pinned to a specific dataset version with recorded parameters, a log and an output dataset.

The notebook is not a record anybody would defend

Scanned pages, editable documents, no attribution, no countersignature. It documents the work but does not evidence it.

Entries move from draft to submitted to signed to witnessed, all attributed. Signing locks the entry, and the signer cannot witness their own work.

Provenance gets reconstructed rather than recorded

Somebody spends two days before a submission rebuilding which run produced which number.

Each run, entry, dataset version and pipeline run records what it came from. The chain already exists as data: it is read, not rebuilt.

The bench stalls for the least interesting reasons

The reagent expired, the sample is in somebody else's freezer box, and two people booked the same instrument for Thursday.

Inventory derives its own stock status on every adjustment. Instrument bookings reject any overlapping time window outright.

Compliance evidence was never collected in the first place

An ethics review or internal audit asks who did what and when, and the system was never asked to keep that.

Every action is attributable on a searchable activity timeline. The evidence accumulates as a by-product of ordinary work.

04 What you get

PROJECTS

A home for a body of research

Status from planning to archived, a principal investigator, field of study and tags. Every experiment belongs to one.

EXPERIMENTS

A designed study with a lifecycle you cannot skip

Hypothesis, linked protocol, structured variables. Transitions are checked, not typed in.

RUNS

Executions that are numbered for you

Every execution gets the next sequential run number automatically, with its own parameters and observations.

NOTEBOOK

A signable, witnessable laboratory record

Entries tie to a project, experiment and run, and progress from draft to signed to witnessed. Signing locks the record.

PROTOCOLS

Methods as shared, versioned capital

Materials, safety notes and ordered typed steps. Publishing freezes the version.

DATASETS

Data with a history that cannot be rewritten

Immutable versions with a checksum, row count and change note. New versions append, never mutate.

ANALYSIS

Pipelines bound to an exact data version

A pipeline run executes over one named dataset version, recording parameters and a log.

INVENTORY

The physical lab in the same workspace

Quantity, location, lot, expiry and a safety data sheet. Stock status is derived on every adjustment.

INSTRUMENTS

Bookable hardware, without the argument

A booking reserves a time window and is rejected outright if it intersects an existing one.

05 Put to work

The same record, read from fourteen different chairs. Each of these is a real path through the product.



RESEARCH SCIENTIST

Design a study once, and never lose the thread again


Each execution starts a run, numbered by the system, carrying its own parameters and observations against the design that produced it.



RESEARCH SCIENTIST

Write up the work in an entry that would survive a review


Signing stamps you and the time and locks the entry. A colleague witnesses it, and the signer cannot witness their own entry.



LAB MANAGER

Publish a method once and let the version do the work


Publishing freezes the version, so a result always cites the exact edition of the method that produced it.



DATA STEWARD

Version the data so a figure can name its source

Every result names a version that still exists in exactly the state it was consumed in, checksum included.



COMPUTATIONAL RESEARCHER

Re-run the analysis and get the same answer

A result tied to a real execution over an exact data version, which is what makes it something a reviewer can check rather than accept.



LAB TECHNICIAN

Stop losing a week to an empty bottle

Stock status that maintains itself as the bench is worked, rather than a list that was accurate the last time somebody updated it.

Core facility manager. Double-booking becomes structurally impossible.

Research auditor. A review that produces evidence without write access.


Research engineer. Automation under identical governance to the application.

06 Built AI-native


AI that can point to the protocol, the run and the data behind what it says

A suggestion that cannot name the method version and the runs it reasoned over is not a weaker feature, it is a liability, because it launders a guess into apparent authority.

So the accountability was shipped before the autonomy: one route for every AI call, a short approved model list your lab can pin, permission checks identical to the rest of the product, and an append-only record of every call that no role can edit.




WHAT THE AI IS BUILT ON




Something real to cite, already in place

Frozen protocol versions, immutable checksummed dataset versions, system-assigned run numbers, entries with a signer and a separate witness.



Every AI call is logged, and stays logged

One entry per call: workspace, person, trigger, model, summary, records referred to, tokens, latency, cost and outcome. Failures exactly as successes.



Your lab picks the model

Seven approved choices across four providers. Pin the one your lab validated; a model outside the catalogue is refused rather than accepted and hoped for.



The AI inherits your permissions

No AI service account with wider reach. An identifier guessed from another lab comes back as not found rather than forbidden, revealing nothing.

SHIPPED

A research assistant on every screen

Available throughout the application rather than parked on a separate page, wired to your research records.

45 agent tools across the whole lab

Not a read-only demonstration: an agent can create, update and transition real records over Model Context Protocol.

It knows which record you have open

Eight kinds of research record recognised from the page you are on. You never paste an identifier.

Irreversible actions ask first

A first attempt is refused with a reason, and the agent is told the action did not happen.

No subject beats the wrong subject

On a list view the assistant is given no entity rather than guessing at whichever one happens to be first.

Nobody can edit the AI record

No update and no delete operation exists anywhere in the schema. Corrections are new entries that reference the one they supersede.

DESIGNED, TICKETED, NOT BUILT

Nothing in the product calls a model yet. The decision engine below is named as what it is.

Protocol drafting into your library

A research aim and its constraints produce a draft protocol with ordered steps written into the real library, not prose in a chat window. A human publishes it, at which point it freezes.

A recommended next run

Reads prior runs, their results and the notebook entries around them and proposes the next condition or parameter change. Recorded for a researcher to accept or override, never enacted by itself.

Cited search over your own corpus

Searches your published protocols, notebook entries, experiments and dataset versions and returns the exact matched version as a resolvable link.

A scheduled operations loop

Reviews stalled runs, incomplete pipeline runs, quiet projects and reagents below reorder level on a cadence. Behind a per-workspace off switch and spend caps.

AND WHAT KEEPS IT IN BOUNDS

One route for every call

Per-request choice first, then the workspace's pinned model, then the platform default. The model is a setting rather than a surprise.

Sensitive content is refused, not stored

Summaries are scanned before insert for identifiers that have no business being there, and a match is rejected.

A pinned model is called, not routed around

If it is unavailable you get a clear failure instead of a silent switch to a model your lab never validated.

Signed records stay outside AI reach

A signed or witnessed notebook entry is never mutable by an AI path, directly or by cascade. That guarantee predates the AI work.

The most senior role is read-only here

The principal investigator is administrator over every other record type and holds view access only on the AI record.

Drift in the evaluation set fails the build

34 cases across eight categories, including prompt injection and destructive intent, checked against the real schema on every change.

THE AI LAYER IN NUMBERS

45

Agent tools across the research domain and health

6

Actions annotated destructive and gated behind explicit confirmation

8

Research record types recognised from the page you have open

7/4

Approved models across four providers, with one marked default

0

Update or delete operations on the AI record anywhere in the schema

0

Rows in the AI record today, because nothing in the product calls a model yet

STATED PLAINLY

LabsOfScience has shipped the accountability and not yet the autonomy.

Live: the assistant with page context, credit metering, 45 governed agent tools with confirmation on the six irreversible ones, four provider connections, a seven-model catalogue, and an append-only record that no role can edit.

Not live: any code path in the product that calls a model. We would rather show you an empty trail you can verify than a full one you cannot.

07 Who does what

Six research roles are seeded into every workspace automatically, alongside organisation and billing presets.

Principal investigator

Owns the lab's permanent research record, and is the only role that can archive one. Views the AI record, never edits it.

Lab manager

Runs the physical lab — instruments, bookings, stock, safety documentation — and deliberately cannot rewrite the science it supports.

Research scientist

The day-to-day bench and dry-lab role, with reach everywhere it works and archive rights nowhere.

Data steward

Curates the data-provenance layer — datasets, versions and pipelines — without touching the bench science that produced it.

Lab technician

Executes methods rather than authoring them, and records what actually happened: runs, observations, deviations, stock and bookings.

Research auditor

Reviews the complete record end to end and changes nothing, by construction. Cannot sign, witness, transition or archive.

08 Under the hood

A federated GraphQL service on a multi-tenant research platform. Every request goes through an API gateway that validates identity and enforces authorisation at the edge, so the product cannot be reached around it.

The domain is modelled properly rather than as documents with fields. Experiments and pipeline runs move through validated state machines. Run numbers are assigned server-side, dataset versions append and never mutate, publishing a protocol freezes its version.

File bytes are never stored by the research service. Around the domain sits the shared platform: identity, access, tagging, activity, billing, notifications and accessibility, inherited rather than rebuilt.

THE SPLIT THAT MATTERS

The product owns the research domain. It owns none of the cross-cutting concerns — identity, permissions, files, comments, tags, audit, billing, notifications and localisation come from shared services. That split is why the product ships enterprise-ready, and why a security control fixed once applies everywhere.

CAPABILITY BY AREA

Authorisation at the edge

Every query and mutation carries a directive the gateway enforces before the request reaches the product. Eleven record types, no unguarded operation.

Reproducibility primitives

Server-assigned run numbers, immutable checksummed dataset versions, publish-time protocol freezing, pipeline runs bound to one named version.

Tenant isolation on every write

Mutations owner-filter by workspace and assert a single affected row, so a valid identifier from a different workspace fails as not found.

Record integrity in the notebook

Signing locks an entry. Only the author may sign it, and the signer is rejected as their own witness — in service logic no permission level can relax.

Validated lifecycles

Status is a state machine the service enforces, not a string the caller sets.

Retention and soft delete

A retention worker hard-deletes soft-deleted rows across all thirteen models once the window has passed, so archived does not silently mean kept forever.

OPEN STANDARDS IT SPEAKS

GraphQL, as a federated schema composed with the wider platform

OpenID Connect single sign-on for web and application clients

OAuth 2.0 device-code flow for terminal and desktop clients

OAuth 2.0 client credentials for service-to-service integrations

Model Context Protocol, for AI agents and assistants

Cursor-style pagination on every list endpoint

Prometheus-format metrics, webhooks for outbound events

SCIM user provisioning, available on enterprise agreements

HOW IT DEPLOYS

Multi-tenant cloud

The default. Your workspace is isolated by construction, with no infrastructure for the lab to run, patch or back up.

Dedicated tenancy or self-hosting

Enterprise options, scoped per engagement rather than downloadable.

Regional hosting

Research data is hosted in India today. Additional residency options are an enterprise discussion.

Backups and release process

Point-in-time recovery over thirty-five days, staged rollout, and a human approval before anything reaches production.

09 Security & trust

Research data carries obligations ordinary business data does not. These are the controls the product enforces today.

Authorisation on every operation

Enforced by the gateway before the request reaches the product. One place where access is decided.

Tenant-scoped writes

The response says not found rather than forbidden, so it reveals nothing.

Least-privilege role presets

Six roles, each holding a level per record type, with omission as the denial.

Two-person record integrity

Only an entry's author may sign it, and the signer is rejected as their own witness.

Immutability where it matters

Signed entries lock, dataset versions never mutate, published protocols freeze.

Encryption in transit and at rest

Secrets held in a managed vault rather than in configuration.

Full activity trail

Searchable and exportable, filterable by actor, action or date for a review.

Sensitive-content screening on write

Freeform write paths reject content carrying identifiers that should not be stored.

Supply-chain and build controls

Vulnerability and secret scanning on every build, and a human approval step before production.

STATED PLAINLY

Nothing here is certified yet, and it would be dishonest to imply otherwise.

SOC 2 TYPE II

Observation period scheduled to begin October 2026, report targeted for Q2 2027.

ISO 27001

Certification targeted for Q3 2027, with a privacy-management extension following it.

21 CFR PART 11

No validation performed. Gap assessment scheduled for December 2026. HIPAA is not supported by default.

If your procurement process requires a current attestation, we are not the right answer yet, and we would rather you knew that before the pilot than during it.

10 One product, many front doors

One record, reached from wherever the work is happening. Everything else is named as what it is.

AVAILABLE

Research web application
Federated research API
Model Context Protocol server, 45 tools
Research assistant
Application shell

IN PROGRESS

Documentation site
Mobile applications for iOS and Android, scaffolded for bench-side capture. Store publishing has not happened.

PLANNED

Desktop application
Command-line interface
Node and Python SDKs
Editor and browser extensions
Community hub



11 Plans

Published as annual seat bands with two rate cards, academic and commercial, because grant cycles are annual.

FREE

\$0

3 seats. An individual researcher or small group evaluating the platform.

The full core loop, not a walled sample.
Community support.

LAB

**\$1,250/yr
academic ·
\$2,500/yr
commercial**

5 seats. A single working lab or research group.

Signable notebook workflow, immutable dataset versions, 45 governed agent tools.
Email support.

GROUP

**\$3,300/yr
academic ·
\$6,600/yr
commercial**

15 seats. A team that has outgrown a single bench.

Shared protocol library, higher capacity, deeper analytics history.
Priority support.

DEPARTMENT

**\$9,000/yr
academic ·
\$18,000/yr
commercial**

50 seats. A department or multi-lab programme.

Multiple isolated workspaces, organisation role presets, single sign-on.

INSTITUTION

**From \$15,000/yr
academic · from
\$30,000/yr
commercial**

100 seats, contracted above. An institution, research office or R&D division.

Custom quotas, dedicated tenancy or self-hosting, grants and billing role presets.

WHAT THE PRICES DO AND DO NOT MEAN

Pre-launch pricing, stated honestly

Only the Free plan exists in live billing today; no paid plan has been created in any environment, and the seat metric these bands depend on is still being built. Treat the numbers as the intended launch configuration.

Why two rate cards

Academic eligibility follows the market-standard definition: academic, government and non-profit organisations. Commercial is set at twice academic. It replaces a discount code with a published price.

Annual only, banded on seats

Billing is never based on how much science you record. A price that taxes the core loop would push a rational team into batching or skipping records, which is precisely the provenance chain the product exists to protect.

What sits outside the plan

AI model usage meters against workspace credits. Production single sign-on, audit streaming, data residency, multi-entity structures, dedicated tenancy and self-hosting are enterprise options.

12 Why teams choose it

Reproducibility is a property of the system, not a discipline

No tracked competitor ships the full set, and the ones that ship parts of it mostly rely on the user remembering.

The physical and digital lab share a workspace

A reagent lot and the instrument time behind a result are part of the same record, not a separate operational tool.

One data model, not integrated modules

Every competitor started as a point solution and grew outwards, so the seams show and the data drifts.

Published prices with two rate cards

Against a category where the larger vendors route everything through a sales conversation and a custom contract value.

Provenance is captured, not curated

Nobody assembles the chain by hand before a submission, which is why it is still correct a year later.

The API is the product, not an export

Automation, agents and the user interface cannot diverge in what they are allowed to do.

HONEST ANSWERS TO HARD QUESTIONS

You are pre-launch. Why would we bet a lab on you?

You should weigh that seriously. There are no customers, no revenue and no reference labs. What does exist is a research system of record across four domains, a deployed web application, 45 agent tools, and a backend end-to-end suite run against production. The honest framing is early-adopter.

We are a regulated lab. Can we use this for good laboratory practice work?

Not as a validated system, and we will not pretend otherwise. Signatures, witnessing and audit trails are real product features, but no 21 CFR Part 11 validation has been performed. If your quality system requires a validated vendor now, we are not that vendor yet.

Migrating a lab off its current notebook sounds like a year of pain.

Migration tooling is a genuine gap. There are no vendor-specific importers today; the practical route is the programmable interface plus an implementation engagement. Most labs start with a new programme rather than backfilling a decade.

Does our research data train anybody's model?

Every AI read and write is scoped to your workspace and permission-checked, and a lookup across labs returns not found. The AI record stores short non-identifying summaries and references, never raw prompts.

WHERE IT SITS AGAINST THE ALTERNATIVES

VS ENTERPRISE R&D SUITES

We do not try to out-build them in biologics tooling. We offer one connected lifecycle rather than assembled modules, and a published price where they require a sales conversation.

VS ACADEMIC NOTEBOOKS

For a lab that only needs a notebook they are a sound choice. We differ in that inventory and scheduling sit in the same workspace, runs are numbered and datasets versioned.

VS OPEN-SOURCE NOTEBOOKS

Zero licence cost and full sovereignty is a real advantage we cannot match. We are not open source. We add dataset versioning, reproducible pipelines and structural provenance.

VS INVENTORY TOOLS

We have no procurement layer at all and intend to integrate rather than rebuild one. We add the research context: a reagent lot that traces to the run and the result it supported.

VS PROTOCOL REPOSITORIES

They own public method discoverability, a role we would rather export to than replace. We bind a method to the experiments that ran it, inside workspace permissions.

VS SPREADSHEETS AND DRIVES

The real incumbent, and it deserves credit: free, universally understood, infinitely flexible. It fails at one thing, which is holding a record together over time and across people.

START TODAY

See whether it holds up against your own lab

The most useful next step is not a demonstration of a happy path. It is bringing a study you already ran and checking whether the record survives the trip.

Join the waitlist for early access and founding-customer terms ahead of the 2026 launch.

Book a walkthrough tailored to your discipline, your instruments and your compliance posture.

Bring one real experiment and see whether the chain holds end to end.

JOIN THE WAITLIST

labsofscience.com

BOOK A WALKTHROUGH

labsofscience.com/contact

INSTITUTION AND ENTERPRISE

sales@labsofscience.com

SECURITY AND DISCLOSURE

security@labsofscience.com
