

ENVIRONMENT · SUSTAINABILITY · CONSERVATION

EcolImpactHub

One workspace for carbon accounting, environmental monitoring, biodiversity and ESG reporting — where every figure keeps the record it came from.

CARBON LEDGER

SITE MONITORING

BIODIVERSITY

ESG REPORTING

AI-NATIVE

THREE AREAS, ONE RECORD MODEL

Measurement and monitoring
Wildlife and conservation
Sustainability and reporting

One workspace, one site reference, one audit trail.

01 The short version

EcoImpactHub is one place to measure your environmental impact, look after the habitats and species you are responsible for, and prove both to a regulator, a board or a donor.

EcoImpactHub replaces that stack with one workspace built on a single record model. A water-quality reading, a camera-trap sighting, a fuel-consumption entry and a reforestation plot all live in the same place, anchored to the same site, under the same identity, with the same evidence trail attached. Three areas of the product — measurement and monitoring, wildlife and conservation, sustainability and reporting — share that model rather than sitting beside each other.

The consequence is that reporting stops being a separate annual project. A disclosure assembles from data you already captured, and each disclosed figure keeps a link back to the reading, sighting or ledger entry beneath it. The evidence comes first, and the report is what falls out of measuring well.

THE PROMISE

Impact rests on evidence, never spin. If a figure cannot name where it came from, it does not belong in the report.

In numbers

3

Areas, one record model

Measurement and monitoring, wildlife and conservation, sustainability and reporting — sharing one workspace, one site reference and one audit trail.

8

Reporting frameworks

GHG Protocol, GRI, TCFD, CDP, SASB, CSRD, ISSB and EU Taxonomy, plus a custom report builder for anything else you answer to.

40

Tools for AI agents

The Model Context Protocol surface exposes forty domain tools, so an outside agent works through exactly the permissions you hold.

Day 1

Governance included

Role-based access, an audit trail, multi-tenant isolation and quotas arrive with the first workspace, not in a later enterprise upgrade.

Who it is for

Sustainability and ESG leads

You own the footprint and the disclosure. You need a carbon ledger that survives assurance, a baseline you can reduce against, and framework-mapped reports whose line items point back at real activity rather than a reconciled guess.

Government and municipal environment teams

You monitor air, water and green space, carry statutory obligations, and report on environmental performance. Sites, sensors, thresholds and framework reports give you a defensible record instead of a quarterly assembly job.

Conservation NGOs and protected-area teams

You run monitoring and restoration on constrained budgets and have to prove impact to the funder who paid for it. A species catalogue, method-tagged sightings, project budgets and a dated activity log turn a season of fieldwork into something you can show.

Field rangers, biologists and restoration crews

You are the largest group of logins and the source of most of the data. Capture is structured and quick, you can only add records rather than reshape catalogues, and your sightings of protected species are location-protected by default.

Environmental consultancies

You design methodologies and run measurement programmes across many clients at once. Every client programme sits in its own isolated workspace with role-scoped access, so one platform serves the whole book without data ever mixing.

Assurance providers, compliance officers and funders

You have to trace a published figure back to the record that produced it, and you must not be able to change anything. A read-only role spans the evidence base with no write path anywhere in it.

Research agencies and scientists

You need rigorous, structured capture — species, method, count, coordinates, timestamp, evidence — rather than free text, and a record model that keeps provenance on every observation so the dataset holds up years later.

What breaks today

None of these are exotic problems. They are the ordinary, structural failures of running environmental work on tools that were never designed to meet each other, and they show up in the same order in almost every organisation.

The data lives in five tools that never speak

The carbon spreadsheet, the mapping layer, the monitoring database, the field notebook and the ESG deck each hold a slice of the truth. Reconciling them is a recurring job that produces no new insight, only agreement.

- Sites, sensors, readings, species, sightings, restoration projects, carbon entries, targets and reports share one workspace and one site reference, so the areas inform each other instead of drifting apart.

The number in the report cannot be traced back

An auditor asks where a figure came from and the answer is a folder, a colleague on leave, or a spreadsheet tab nobody wants to open. The claim then survives on trust rather than evidence.

- Every reading records how it was captured, every carbon entry names its factor and scope, and every report line item points at the record beneath it — so a reviewer following a figure lands on a measurement.

Reporting is an annual scramble

Disclosure becomes a separate project each year: chase a dozen owners, hope the numbers agree, publish, and quietly lose the working papers before the next cycle starts.

- A report is scoped to a framework and a period and assembles from line items drawn from data already captured, then moves through draft, review and published — so disclosure is continuous rather than annual.

A silent sensor looks exactly like a healthy environment

A flat line reads as calm water until somebody notices the probe stopped reporting a fortnight ago. The most dangerous reading is the one that never fires.

- Every device carries an explicit health state — online, offline, degraded, in maintenance or retired — and a last-reading time, so a gap in the data is a visible fact instead of a quiet assumption.

What breaks today, continued

Conservation outcomes stay anecdotes

Plantings, invasive removals and surveys are held on paper and in memory, so the arc from bare ground to recovering habitat cannot be shown to the funder who paid for it, or checked by the team who did it.

- A restoration project carries its boundary, target area, budget and dates, and each intervention is logged as a dated, quantified activity as the project moves through its lifecycle. The record is the report.

Field data arrives as a transcription backlog

A day in a remote reserve becomes a pile of paper forms somebody has to key in later, with the inevitable losses: illegible counts, missing coordinates, dates reconstructed from memory.

- Observations are captured against a known species with a detection method, a count, a timestamp, coordinates and optional photo or audio evidence, so records from different observers and sites line up cleanly the moment they land.

Protecting a species can be the thing that exposes it

Precise coordinates for an endangered sighting are a poaching risk, not a privacy nicety, yet most tools treat location as ordinary metadata and hand it to anybody with read access.

- Coordinates for endangered and critically endangered sightings, and for priority protected sites, are generalised to roughly a kilometre on every ordinary read, with no setting anywhere to switch it off. Precise locations require a separately granted permission.

Carbon tooling stops at carbon

A platform can tell you how much you emitted and nothing about what it did to the watershed downstream, what lives there now, or whether the restoration you funded is working. Half the environmental picture goes unmeasured.

- The carbon ledger sits beside site monitoring, biodiversity and restoration in the same product, so the footprint and the ground it affects are measured with the same rigour and the same evidence trail.

What you get

SITES

The registry every number hangs off

A monitored place — facility, watershed, forest, wetland, farm, protected area, airshed, coastal zone or other — with a type, a region and a status. Archiving retires a site from the active roster without deleting its history. Sensors, readings, sightings and projects all attach here.

SPECIES

The catalogue and the sightings against it

A workspace-level catalogue entry with a taxon group, an IUCN-style conservation status and a priority flag. Observations record it being detected — by direct sighting, camera trap, acoustic, tracks and signs, environmental DNA, telemetry or community report — with a count, a time, coordinates and optional photo or audio evidence.

TARGETS

Turning a footprint into a plan

A reduction target with a metric, a baseline value and year, and a target value and year, shown against its baseline. It moves through proposed, active, achieved, missed and retired, so progress — or backsliding — is impossible to leave unstated.

SENSORS

The instruments and sampling stations

Air quality, water quality, weather, soil, energy meters, flow, acoustic and camera traps, plus a manual kind for human-collected samples. Each carries a health state and a last-reading time, so the physical network has its own operational view separate from the data it produces.

HABITAT

Restoration projects and their activity log

A coordinated restoration effort with a site, a target area in hectares, a budget, a manager and a lifecycle from planned through in progress and monitoring to completed. Each intervention is logged with a type, a quantity and a date: planting, invasive removal, rewilding, water management, soil remediation, community engagement or survey.

REPORTS

Framework disclosures built from line items

A report scoped to a framework and a reporting period, assembled from line items that each disclose one metric with a label, value and unit. Emissions lines can reference the carbon entry behind them. Reports move through draft, in review and published, and publishing is its own permissioned step.

READINGS

Measurements with provenance and thresholds

A single time-stamped value of one parameter at one site, tagged by source — person, sensor, satellite, import or calculation. Configurable rules fire above, below or outside a range at info, warning or critical severity, and the breach flag is stored on the record so history cannot drift when a rule changes.

CARBON

Emission factors and the ledger

Factors are first-class records carrying a scope, an activity unit, a value per unit, a source attribution and a valid year. An entry captures activity data against a chosen factor, classified by GHG Protocol scope, and moves through draft, recorded and verified so the ledger always distinguishes confirmed from merely entered.

METRICS

One agreed figure per metric per period

A normalised library keyed by metric, category and period, tagged with the source it came from. Entries are upserted by key, so refreshing a period updates the figure in place rather than accumulating duplicates and there is exactly one authoritative number to disclose.

Put to work

The same nine records support very different jobs. Below are sixteen of them, written from the point of view of the person who has to do the work, with the current honest limits noted where they matter.

SUSTAINABILITY LEAD

Build a carbon footprint that survives assurance

Activity data — fuel burned, electricity consumed, business travel, purchased goods — is captured against curated emission factors that carry a source and a valid year. Each entry is classified by GHG Protocol scope and moves through draft, recorded and verified. Totals roll up by site, category and period against a baseline you set.

- A single auditable ledger where the biggest sources are obvious and every figure is traceable to the activity that produced it.

ENVIRONMENTAL PROGRAMME DIRECTOR

One registry for every place you steward

Each location is created once with a name, a type — facility, watershed, forest, wetland, farm, protected area, airshed, coastal or other — and an optional region and country. The catalogue filters by type and status, and archiving retires a site from the active roster while keeping its past readings and projects intact.

- A clean spatial spine, so every downstream view reads cleanly because the places underneath it are curated rather than guessed.

MONITORING ENGINEER

Know which of your instruments are actually awake

Each device is registered against the site it watches with its kind, and carries an explicit health state — online, offline, degraded, maintenance or retired — that you flip inline when a unit goes in for service. Last-reading times stay fresh as readings arrive, and the fleet view filters by kind and by status across every site.

- An operational view of the physical network itself, so a gap is a visible fact and a field visit goes to the right unit.

ENVIRONMENTAL SCIENTIST

Keep watch over water quality, continuously

Readings arrive from manual field samples or connected sensors, each tagged to a site and a timestamp. Parameters such as pH, turbidity and dissolved oxygen build into a time series, and configurable thresholds fire above, below or outside a range at the severity you choose, notifying whoever set the rule.

- A defensible time-series picture of a water body's health, with drift surfacing the moment it happens rather than at the next site visit.

Put to work, continued

FIELD RESEARCHER

Turn a day in the field into queryable data

Every sighting is recorded against a known species with a detection method, a count, the date and time observed, latitude and longitude, an optional site and free-text notes, with photo or audio evidence attached.

Because the fields are standardised, records from different observers line up cleanly.

- Dated, located, queryable records at the end of the day instead of a transcription backlog at the end of the season.

CONSERVATION BIOLOGIST

Answer what lives here, and is it increasing

Species are catalogued once with a taxon group, an IUCN-style conservation status and an optional priority flag. Observations then stream in from surveys, camera traps, acoustic sensors, environmental DNA, telemetry and community reports, each linked to a species, a place and a time, and accumulate against the catalogue entry.

- Scattered sightings become an ecological signal, and the conservation-status snapshot shows where attention is most needed.

WILDLIFE RESEARCHER

Follow a species by its located detections

Each fix is recorded as an observation on the species with a telemetry or direct-sighting method, a count, a timestamp and coordinates. Over time the located records accumulate into a time-ordered detection history on the species page.

- A per-species detection history you can review. Automatic collar ingest, movement paths, home-range polygons and corridor analysis remain on the roadmap.

RESTORATION MANAGER

Document the arc from bare ground to recovery

A project captures its scope — an optional site, a target area in hectares, a budget, start and end dates — then each intervention is logged as a dated activity with a type, a summary, a quantity and a unit. The project advances through planned, in progress, monitoring and completed.

- A documented record rather than a hopeful estimate: the team sees what has actually been done, and the funder sees what their support produced.

Put to work, continued

COMMUNITY COORDINATOR

Fold community reports into the same dataset

A steward's sighting, or an incident such as illegal dumping, is logged as a field observation with the community-report method, against the relevant species, with a count, a time, coordinates or a site, and notes for the context. It sits in exactly the same dataset as professional surveys and camera-trap records.

- More eyes on the ground and richer data than any single team could gather. Dedicated steward enrolment and discussion spaces are still ahead.

ESG MANAGER

Assemble a disclosure you can stand behind

A report is started against the framework you answer to — GHG Protocol, GRI, TCFD, CDP, SASB, CSRD, ISSB, EU Taxonomy or a custom one — for a defined period. A line item per disclosure carries a label, value and unit, and emissions lines can reference the carbon entry beneath them. Publishing is its own permissioned step and stamps a date.

- A structured, traceable artefact rather than a document, so an assurance provider follows the chain instead of taking a number on faith.

COMPLIANCE OFFICER

Make the surprise audit uneventful

Carbon entries are kept verified in the ledger, then a report is assembled against the framework and period you are accountable for, with a quantified line item per required disclosure tracing back to the data. When it is complete you publish it and it carries a publication date.

- A submission becomes an export of an existing defensible record. Dedicated obligation, permit and deadline tracking with reminders is still on the roadmap.

PROGRAMME DIRECTOR

Show donors exactly what their funding changed

The funded restoration project records the budget and target area the funding supports, then the work delivered — plantings, removals, surveys — is logged as dated, quantified activities as the project advances. At the period's end an impact report assembles from the same monitoring data, a line item per outcome, and is published.

- A shareable, evidence-backed record of where the money went and what it changed. Dedicated grant objects and donor portals remain ahead.

Put to work, continued

PROCUREMENT LEAD

Make the value chain visible in your footprint

Scope 3 emission factors are defined for upstream categories — purchased goods, materials, transport — and activity is recorded against them so each entry's emissions are computed. Filtering the ledger to Scope 3 and sorting by emissions surfaces the categories that weigh most, beside Scope 1 and 2.

- An honest view of how much of the footprint sits upstream. Structured supplier mapping and data-request workflows are on the roadmap, and are the clearest gap against Scope 3 specialists.

SUSTAINABILITY COORDINATOR

Turn workshop ambitions into tracked commitments

Each target is set with a metric, a baseline value and year, and a target value and year, and shows a planned-reduction bar against its baseline. An inline status control moves it through proposed, active, achieved, missed and retired, and the ledger activity that will move the number sits alongside it.

- We are working on sustainability becomes a concrete commitment with a baseline, a target year and a status somebody owns.

OPERATIONS MANAGER

Prove the cleaner energy switch is paying off

An emission factor is defined for your grid or supply, then consumption is recorded as Scope 2 ledger entries tagged to a site so emissions are computed for each entry. A reduction target with a baseline and a lower target year makes the intent trackable, and Scope 2 appears in the carbon-by-scope view.

- Energy shows up directly in the footprint. Live generation monitoring, energy-mix dashboards and automatic avoided-emissions calculations are on the roadmap.

PLATFORM ENGINEER

Let the AI agent you already use drive the product

A Model Context Protocol server exposes forty domain tools spanning sites, sensors, thresholds, readings, species, observations, habitat projects, restoration activities, emission factors, carbon entries, targets and reports, plus four read-only resources and two guided workflows. Every call goes through the same permission checks as the interface.

- An agent does real work under your permissions, with no privileged back door and no access to precise protected-species coordinates.

06 Built AI-native

AI that can show you the reading behind every number

An environmental number is only worth what it can prove. An auditor, a carbon buyer or a ranger will always ask where it came from, which makes a confident but unsourced answer worse than no answer at all. So the order of construction matters: EcoImpactHub records the evidence first — every reading names how it was captured, every carbon entry names its factor and scope, every report line names the figures beneath it — and the AI works on top of that rather than beside it.

That has three practical consequences. Answers are built from your live records rather than a separate copied index, so a fact carries a source and a timestamp. Carbon, monitoring and wildlife sit in one record model, so a model can reason across all of them in a single pass instead of stitching four exports together. And the AI inherits your permissions rather than holding its own, including the coordinate generalisation that protects endangered-species sightings — there is no separate AI route into the data.

The second half of that discipline is refusing to overstate what is built. The assistant, the agent surface, per-turn cost metering, threshold evaluation, location protection and the audit trail are live. The domain judgement — anomaly triage, emission-factor matching, cited report drafting, species identification, the overnight sweep — is designed against live data and none of it runs today. Each is described below with its real status, because a catalogue that blurs the two is exactly the kind of document this product exists to replace.

What AI-native means here

Every number keeps its receipt

Each reading names how it was captured, each carbon entry its factor and scope, each report line the figures beneath it. An answer points at a source rather than asserting one.

Carbon, monitoring and wildlife in one place

Sites, sensors, readings, species, sightings, restoration, carbon and reports sit together, so AI can reason across all of them in one pass instead of reconciling four disconnected tools.

Your permissions, not the model's

The assistant and any agent see exactly what you can, never more. Sensitive wildlife locations are generalised before anything reads them, AI included, with no switch to disable it.

Agents use the same product you do

There is no separate AI back door. Outside agents register sites, log readings and assemble reports under the same permissions, the same limits and the same audit trail as the interface.

What the AI does today

Assistant on every screen

Live. A chat panel opens from the dashboard, sites, monitoring, species, observations, habitat, carbon and reports, so it follows you rather than living on one page.

Each conversation runs inside your own workspace, private to it, carrying your own permissions.

Sources recorded on the way in

Live. A reading records whether it came from a person, a sensor, satellite, an import or a calculation, plus its site, parameter and time.

You choose the source as you log the data, so anyone — or anything — reading it later can trace it back.

Forty domain tools for outside agents

Live. Sites, species, sensors, thresholds, readings, observations, restoration, carbon, emission factors, targets and reports — enough to stand up a monitored site or assemble a disclosure end to end.

Connect any Model Context Protocol client and the agent signs in as you, working on real records with nothing bespoke to build.

Agent context and guided workflows

Live. Four read-only resources let an agent pull the schema, your sites, your reduction targets and your carbon ledger to orient itself before acting, alongside two scripted walkthroughs.

Ask an agent to set up a new monitoring site and it follows the steps in order rather than guessing at them.

Platform tools in the same agent session

Live. The agent surface also reaches files, tags, conversations, notifications, scheduling, export, organisations and access control.

That is what makes a real workflow possible: attach the field photo, tag the observation, post the summary, schedule the follow-up.

What the AI does today, continued

Threshold breaches flagged as data lands

Live. Each reading is tested against your above, below and outside-range limits at info, warning or critical severity at the moment it is saved.

The flag is stored on the record, so changing a rule later cannot rewrite history, and whoever set the rule is notified.

One agreed figure per metric per period

Live in the backend. The normalised metrics library is keyed by metric and period, source-tagged, and upserted so a refresh updates in place rather than duplicating.

There is a single authoritative number to disclose, which is what makes a cited answer possible. A dedicated interface for it is on the roadmap.

Five moments automations can react to

Live. Site created, reading recorded, threshold breached, carbon entry verified and report published are published as domain events today.

Point an automation at any of the five and it starts the moment that thing happens. Nothing in the product slows down if a listener goes offline.

AI spend counted turn by turn

Live. The assistant checks your remaining allowance before a turn and reports what the turn consumed, along with which provider and model answered.

AI cost is a line you can watch rather than an end-of-month surprise. The workspace allowance is the hard limit; the check beforehand is an early warning.

What is not built yet

Grounding in the record you are viewing

In progress. The assistant knows the product but not what is on your screen, so it cannot yet answer which of these entries is still unverified.

Once it lands, opening a site and asking a question returns an answer about that site, with no references pasted into the chat.

Species suggestions from field evidence

Roadmap. Planned to read the photo or audio attached to an observation and suggest a species from your own catalogue, returning unknown rather than forcing a match.

Built to triage a six-week camera-trap backlog down to the frames a specialist genuinely needs to see, with sensitive locations still generalised.

Failing sensor versus real event

Roadmap. Planned to catch slow drift that never crosses a limit, step changes aligned with maintenance windows, and physically implausible values.

It will cite the readings and the sensor's own health history behind every call, so a monitoring engineer can look and disagree.

An overnight sweep that escalates

Roadmap. Planned to check nightly for sensors that missed their reporting window, unacknowledged breaches, draft carbon entries near a deadline and dormant restoration projects.

It will read widely and write almost nothing — verifying, publishing and archiving stay human decisions.

Suggested emission factor and scope

Roadmap. Planned to propose the closest factor from your own library rather than an external black box, propose the GHG Protocol scope, and pre-fill a draft entry.

No confident match will be a valid answer. You still record the entry and verify it as two separate actions.

Report drafts that must cite evidence

Roadmap. Planned to work a framework's expected disclosures against what your workspace holds and draft the lines under one hard rule.

A line that cannot name the readings or entries behind it is never written. I could not substantiate this is a valid result; an unsourced figure is not.

What keeps it in bounds

The AI inherits your permissions

Every read and write is permission-checked before the product answers, so an assistant turn or an agent call reaches exactly what you can reach.

There is no separate AI route in and no elevated service identity behind the chat panel.

Wildlife locations generalised by default

Coordinates for endangered and critically endangered sightings, and for priority protected sites, are rounded to roughly a kilometre on every ordinary read path, with no setting anywhere to turn it off. This is a poaching control, not a privacy nicety.

Precise coordinates are unreachable by agents

The separately permissioned queries that return exact locations are deliberately not wired into the agent surface at all, so an agent cannot ask for a poaching-relevant location because the tool to ask does not exist.

People verify and publish, not AI

Verifying a carbon entry and publishing a report are separate permissioned steps that no AI path can take. A disclosure an assurance provider will accept cannot rest on unreviewed model output.

Every AI-reachable change is recorded

Changes, plus fourteen sensitive read operations, land in your audit trail with private values stripped, so what an assistant or an agent did stays attributable long after the session ends.

Writes are metered and quota-bound

Anything the AI writes counts against the same allowances as a click in the interface, so an automated actor's activity is countable rather than invisible.

Tools are labelled read-only or destructive

The tool layer marks each tool so a client can warn a user before an archive or a publish. The confirmation gate is enforced at the shell today; wiring it to the six risky product tools is tracked work, not a claim.

Grounded in live records, not a copy

Nothing is duplicated into a separate search index. Answers come from the same records the interface shows you, which is why an answer can cite one.

The model is reported, not pinned

The assistant runs on a shared service that selects a model per turn and reports back which provider and model answered — that report is how spend is counted. Choosing your own model, or bringing your own key, is not available yet.

A review area, deliberately empty

The place where AI proposals will be accepted or rejected exists in the product and says coming soon, because the judgement capabilities that would fill it have not shipped. Every proposal will have to show its citations before anyone can accept it.

Reachable from, and in numbers

REACHABLE FROM

- A chat panel on every screen — dashboard, sites, monitoring, species, observations, habitat, carbon and reports
- Any Model Context Protocol client, including the assistant in your editor or your own automation stack
- Forty domain tools covering sites, sensors, thresholds, readings, species, observations, habitat projects, restoration activities, emission factors, carbon entries, targets and reports
- Four read-only agent resources: the schema, your sites, your reduction targets and your carbon ledger
- Two guided prompt workflows: standing up a monitored site end to end, and taking a carbon ledger through to a framework disclosure
- Shared platform tools in the same agent session — files, tags, conversations, notifications, scheduling, export, organisations and access control
- Five domain events any automation can subscribe to, as the trigger surface for scheduled work
- An AI review area in the product, reserved for proposals to accept or reject, marked coming soon until the judgement capabilities ship

IN NUMBERS

- Forty domain tools on the agent surface, covering the contract from registering a site to publishing a report.
- Four read-only agent resources and two guided prompt workflows ship alongside those tools.
- Five domain events are published today: site created, reading recorded, threshold breached, carbon entry verified, report published.
- Every assistant turn runs a credit check beforehand and reports the tokens actually consumed, plus the provider and model that answered.
- Coordinates for endangered and critically endangered sightings, and for priority protected sites, are generalised on every ordinary read path, with no switch to disable it.
- Fourteen sensitive read operations are recorded in the audit trail alongside every change, with private values stripped.
- Verifying a carbon entry and publishing a report are separate permissioned actions that no AI path can perform.
- Every query and mutation carries an access-control directive enforced at the API edge, so an agent call is checked exactly as a click is.

06 The gaps, named plainly

The gaps are worth naming plainly. The assistant does not yet know which record you are looking at, so it can explain how the carbon ledger moves from draft to recorded to verified but cannot tell you that one site's dissolved-oxygen sensor has been silent for nine days; passing it the page context is the work in flight. None of the domain judgement is built — anomaly detection, telling a failing sensor from a real event, emission-factor and scope suggestions, cited report drafting, species identification from photo evidence and the overnight sweep are all designed against live data and none of them run today, and they are sequenced behind shared platform work rather than close to shipping. The review area exists and says coming soon because there is nothing yet to review. There is no evaluation set measuring whether a generated environmental claim is genuinely substantiated, and no defence that treats free-text field notes as data rather than instructions. Choosing your own model or supplying your own key is not available. And the confirmation gate that would make an agent ask before archiving a site or publishing a report is enforced at the shell but not yet attached to the six risky product tools.

Who does what

Access is expressed as a level — admin, editor or viewer — on each kind of record, and a role simply omits the records it should not reach. The omission is the access control. Six product roles are seeded into every workspace automatically, alongside organisation and billing presets.

Program Owner

Owns the programme end to end. Usually the workspace owner or enterprise administrator.

Accountable for the whole environmental programme, from field capture to published disclosure, and decides who else gets in. The only role holding admin on all thirteen kinds of record.

- Register and archive sites
- Curate the species catalogue
- Publish disclosures
- Commission and retire instruments
- Correct the carbon ledger
- Grant and revoke access

Sustainability Lead

The footprint and the disclosure. Answers to the board and to assurance.

Owns carbon accounting, reduction targets and the reports published from them, with read access to the field evidence behind each figure but no ability to reshape the monitoring backbone.

- Curate emission factors
- Set baselines and targets
- Read sites and readings
- Record and verify entries
- Assemble framework reports
- Publish a disclosure

Monitoring Manager

That the network is awake and the thresholds fire correctly.

Builds and maintains the monitoring backbone — sites, instruments, thresholds and the reading record — and sees carbon rollups for their sites without any authority over the ledger.

- Register sites
- Tune threshold severity
- Review breaches
- Commission and retire devices
- Record readings
- Read carbon rollups

Conservation & Restoration Lead

What lives here, whether it is recovering, and whether the restoration is working.

Runs biodiversity and restoration work: the species catalogue, field observations including protected-species locations, habitat projects and their activity logs. No write path into carbon or thresholds.

- Maintain the species catalogue
- Validate field sightings
- Log restoration activities
- Set conservation status and priority
- Run habitat projects
- Read sites and readings

Who does what, continued

Field Observer

Getting the day's work recorded correctly and quickly.

Captures readings, sightings and restoration work on the ground, and views the sites, instruments, species and projects being worked against. Can add records but cannot reshape catalogues or retune rules.

- Record measurements
- Log restoration activities
- View the species catalogue
- Log observations with evidence
- View assigned sites
- View active projects

Assurance Auditor

Whether a published figure traces back to a real record.

Reviews the evidence behind every disclosed figure across monitoring, conservation and carbon accounting, with no write path anywhere. Read-only across twelve of the thirteen record types.

- Trace a report line to its entry
- Read readings and thresholds
- Read emission factors
- Read the carbon ledger
- Read habitat projects
- Change nothing at all

Billing & Organisation Administrator

Plan, capacity and cost across programmes.

Organisation and billing presets seeded alongside the product roles, covering subscriptions, quotas and invoices across the workspaces an organisation runs.

- Manage the subscription
- Add capacity
- Assign organisation roles
- Review quota consumption
- Access invoices
- View programme rollups

Under the hood

EcoImpactHub is API-first. Product state is served over a single federated GraphQL graph, and the web experience is one client of it rather than a privileged path — anything the interface can do is reachable programmatically with scoped keys. Clients never reach the service directly; requests go through an API gateway that authenticates the caller and enforces access control at the edge before the request lands.

The domain is organised into three modules over one record model: monitoring holds sites, instruments, readings and thresholds; conservation holds the species catalogue, observations, habitat projects and restoration activities; sustainability holds emission factors, the carbon ledger, the metrics library, reduction targets and reports with their line items. Every type is namespaced so it composes cleanly into a shared graph alongside the rest of the platform without ever colliding.

Around that sit the shared services the product deliberately does not rebuild: identity and single sign-on, access control, tenancy, billing and quotas, notifications, files for photo and audio evidence, tags, conversations, scheduling, workflow approvals, activity timelines, internationalisation and accessibility. That is why a release ships enterprise-ready rather than accumulating enterprise features later, and why photo evidence on an observation is managed by the same storage as everything else in your workspace.

THE SPLIT THAT MATTERS

The split that matters is between capture and judgement. Capture is deterministic and complete: a reading is evaluated against your thresholds at the moment it is written, and the breach flag is stored on the record itself rather than recomputed later, so changing a rule tomorrow cannot rewrite what happened yesterday. Judgement — deciding whether that breach is a failing instrument or a real environmental event — is deliberately a separate layer, is not built, and when it arrives will always cite the records it reasoned over. Everything in this catalogue sits on one side of that line or the other, and the line is drawn honestly.

Capability by area

Monitoring

Nine instrument kinds including a manual sampling station, explicit device health states with last-reading times, readings tagged by source, and threshold rules that fire above, below or outside a range at three severities, evaluated at write time.

Conservation

A workspace species catalogue with taxon group, conservation status and priority flag; observations by seven detection methods with counts, timestamps, coordinates and file-backed photo or audio evidence; habitat projects with target hectares, budget and a four-stage lifecycle; seven kinds of logged restoration activity.

Sustainability

Source-attributed emission factors with an activity unit, a value per unit and a valid year; scope-classified carbon entries with a draft, recorded and verified lifecycle; scope totals by site, category and period; baseline-and-target reduction goals; framework reports assembled from traceable line items.

Evidence and provenance

Every reading carries its capture source; every carbon entry names its factor and scope; every report line points at the record beneath it; the metrics library holds one upsert-keyed figure per metric per period, tagged with the source it came from.

Access control

Every query and mutation carries an access-control directive that the gateway enforces from the composed schema before the request reaches the service. Levels are admin, editor and viewer per record type; a role denies a record type by omitting it. Field observations are read-suppressed by default.

Multi-tenancy

Every record is scoped to its workspace, and writes cannot cross a tenant boundary. Consultancies and NGOs run isolated client programmes side by side from one platform; dedicated tenancy is available on Enterprise terms.

Events and automation

Five domain events are published on the platform event stream — site created, reading recorded, threshold breached, carbon entry verified, report published — so downstream consumers react to environmental facts as they happen rather than polling, and nothing blocks if a listener is offline.

Data lifecycle

Every underlying record carries soft-delete and audit columns, so archiving a site or a species retires it from the active roster while its history, readings and projects stay intact and reviewable.

08 Standards and integrations

OPEN STANDARDS IT SPEAKS

- GraphQL, served as one federated graph across the whole platform
- OpenID Connect single sign-on
- OAuth 2.0 — device code for terminal clients, client credentials for machine integrations
- Model Context Protocol, for AI agents and assistants
- GHG Protocol Scope 1, 2 and 3 classification
- GRI, TCFD, CDP and SASB disclosure frameworks
- CSRD, ISSB and EU Taxonomy
- IUCN-style conservation status categories on every catalogued species
- WGS84 latitude and longitude on every located record, with timestamps on every reading
- WCAG 2.1 AA as the accessibility conformance target, with a design-system audit still pending

PLUGS INTO

- Shared evidence storage for observation photos, audio and report attachments, with signed-link access
- Workspace-wide tagging shared across sites, species, projects and other products
- Inline comments and discussion threads on sites, species, projects and reports
- In-app, email and webhook notifications with per-user preferences
- Calendar and scheduling for monitoring rounds, survey timetables and recurring imports
- Multi-step approval workflows for report sign-off, target approval and restoration milestones
- Single sign-on and directory-backed identity across every client
- Subscriptions, regional pricing, usage quotas and invoices
- A searchable, exportable activity timeline recording actor, target and result for every action
- A developer portal issuing API tokens, OAuth applications and scoped keys

How it deploys

Multi-tenant cloud

The default. Strict per-workspace isolation on shared infrastructure, with nothing for you to run or patch.

Dedicated tenancy

Available on Enterprise terms for organisations that need their environmental data on separate infrastructure.

Self-hosting

Offered on Enterprise terms. Scope, support and update cadence are agreed per contract; this is not a downloadable distribution today.

India data residency

All state stays in India by default. That is deliberate, and it means no cross-region deployment and no geo-redundant disaster recovery at present.

API-first automation

Every capability is reachable over the graph with scoped keys, so the product fits inside an existing stack rather than sitting beside it.

09 Security & trust

Security and governance are inherited from the shared platform, so they apply uniformly rather than being reimplemented per product. What follows is what is actually enforced, and the section after it states what is not yet certified.

Authentication at the edge, never in the service

The API gateway authenticates every caller and forwards a signed identity. The product service never validates tokens itself, which removes an entire class of inconsistent-auth defect between surfaces.

Access control on every operation

Every query and mutation carries an access-control directive that the gateway enforces from the composed schema before the request reaches the service. A field with no matching permission is denied — the default is closed, not open.

Multi-tenant isolation by default

Every record is scoped to its workspace and writes cannot cross a tenant boundary. Dedicated tenancy and self-hosting are available on Enterprise terms for organisations that require physical separation.

Sensitive-location protection for wildlife

Coordinates for endangered and critically endangered sightings, and for priority protected sites, are generalised to roughly a kilometre on every ordinary read path. There is no configuration option to disable it, and the agent surface cannot request the precise values at all.

Field observations are read-suppressed

A sighting carries an observer identity together with precise coordinates and a timestamp, so a blanket read-only role receives no sightings by default. Only the programme owner, the conservation lead and the field observer reach them; giving an external reviewer sighting-level detail requires an explicit individual grant.

Separation of recording from publishing

Verifying a carbon entry and publishing a report are their own permissioned actions. No automated or AI path can perform either, because a figure an assurance provider will accept cannot rest on unreviewed output.

End-to-end audit and activity trail

Every mutation is recorded with actor, target and result, alongside fourteen sensitive read operations, with private values stripped. The timeline is searchable and exportable for compliance review.

Security & trust, continued

Multi-organisation governance

Partners, auditors and community contributors collaborate in shared workspaces, each scoped by role, with every action attributable to the individual who performed it.

Payment data never reaches us

Card payments are handled by the payment processors; only tokenised payment-method references are stored, so cardholder data is outside our environment entirely.

Quota enforcement at the action layer

Capacity limits are checked when the action is attempted rather than reconciled afterwards, so there are no retrospective surprises. Five of the nine published quotas are catalogue-only today and are not yet gated in-product.

Continuous dependency and image scanning

Vulnerability, infrastructure-as-code and software-bill-of-materials scanning run on every build, alongside secret scanning and registry image scanning, with findings triaged before release.

Encryption in transit and at rest

All traffic between your clients and the service is encrypted with current transport security, and stored data is encrypted at rest on enterprise cloud infrastructure whose practices are modelled on the ISO 27001 framework.

Published subprocessors

The infrastructure, communication, geospatial and analytics providers that help run the service are disclosed publicly, with launch-critical providers still being finalised marked as such rather than omitted.

Secrets never live in source or client builds

Every environment reads its own secret store, which is the source of truth; configuration files are generated from it and no secret is baked into a shipped client bundle.

Stated plainly

EcoImpactHub holds no security or privacy certification today, and the roadmap should be read as intent rather than assurance. SOC 2 Type II and ISO 27001 are planned with no auditor yet engaged; ISO 27701, 27017 and 27018 sit behind them. GDPR, UK GDPR and India's DPDP Act readiness are in progress — the data processing addendum and privacy policy are published, technical automation of data-subject rights is only partly implemented, and an EU/UK representative will be appointed once the threshold is reached. The payment-card assessment path is narrow but still to be confirmed with the processors. WCAG 2.1 AA is a conformance target with the design-system audit pending, not a certified result. Data stays in India by default, which means no cross-region deployment and no geo-redundant disaster recovery today. The internal security gate passed with accepted risks after mass-assignment, workspace-context and cross-reference access issues were found and fixed; stricter validation of date ranges and string lengths is planned hardening rather than shipped. The backend end-to-end suite exists but is not yet wired into continuous integration, so domain flows are not verified on every deploy. And to be unambiguous: EcoImpactHub is not a certification body or a regulatory authority. It holds the traceable evidence beneath an environmental claim; it does not certify the claim, issue credits, or make a compliance decision on your behalf.

One product, many front doors

The web experience is the complete product today. The rest of the surfaces are specified and sequenced, and each is listed below with its real status rather than an aspirational one.

Web application

The full product — sites, instruments, readings, thresholds, the species catalogue and observations, habitat projects, the carbon ledger, targets, reports and the unified dashboard. Available today.

Product surface in the workspace shell

The same experience composed into the wider workspace application, so environmental work sits beside the other products a team uses under one identity. Available today.

AI assistant

A chat panel reachable from every section, running inside your workspace with your permissions and metered per turn. Live, and honestly not yet aware of the record on your screen.

Agent surface for AI clients

A Model Context Protocol server exposing forty domain tools, four read-only resources and two guided workflows to any compatible client. Built and working; public packaging and distribution are not yet open.

Documentation site

Product guides, API reference and integration tutorials. In progress ahead of launch.

Command-line interface

Terminal-native automation, bulk import and export, and scriptable operations against the environmental API. Specified, install path defined, not yet published.

Front doors, continued

Node and Python SDKs

Typed programmatic access for services and integrations, and an idiomatic Python client for data-science and remote-sensing workflows. Specified, not yet published.

Mobile and field application

Offline-capable field capture with location and photo evidence, plus on-the-go monitoring and approvals. On the roadmap; today field capture happens in the responsive web application.

Browser and editor extensions

Quick capture of evidence from the web, and project resources and API workflows inside your editor. Planned; neither is published.

Webhooks

Signed outbound events for observations, approvals, report milestones and project changes, built on the five domain events already published. Planned.

Community hub

Question-and-answer, discussion and practice sharing, signed in against your workspace identity. Planned for launch.

Professional services

Strategy and consulting, implementation and programme build, data migration, integrations and pipelines, monitoring enablement, analytics enablement, managed environmental operations, governance and assurance, and training. Available and quoted per engagement.

11 Plans

Four self-serve tiers with published prices, and an enterprise shape contracted per organisation. The platform underneath is identical on every tier — what changes is capacity, reporting depth, governance options and support.

FREE

\$0 / month

Individuals, small environmental teams and NGOs getting started with real data rather than a demo.

- 1 monitored site and 10 connected instruments
- 100 measurements a month
- Up to 15 tracked species and 50 observations a month
- Up to 2 conservation programmes
- 1 carbon ledger entry and 2 sustainability reports a month
- 1 restoration hectare tracked
- The same access control, audit trail and tenant isolation as every paid tier

STARTER

\$99 / month

A small programme running its first genuine monitoring and carbon workflow across a handful of locations.

- 3 monitored sites and 100 connected instruments
- 1,000 measurements a month
- 50 tracked species and 500 observations a month
- 10 conservation programmes
- 4 carbon ledger entries and 25 sustainability reports a month
- 10 restoration hectares tracked
- Full framework reporting and evidence-linked line items

GROWTH

\$399 / month

Growing organisations and conservation programmes running continuous monitoring across a real portfolio.

- 15 monitored sites and 1,000 connected instruments
- 10,000 measurements a month
- 250 tracked species and 5,000 observations a month
- 50 conservation programmes
- 12 carbon ledger entries and 500 sustainability reports a month
- 100 restoration hectares tracked
- Priority support

BUSINESS

\$999 / month

Multi-site enterprises and consultancies running isolated client programmes from one platform.

- Unlimited monitored sites under fair use
- 10,000 connected instruments and 100,000 measurements a month
- 1,000 tracked species and 50,000 observations a month
- 500 conservation programmes
- Unlimited carbon ledger entries under fair use, 5,000 reports a month
- 1,000 restoration hectares tracked
- Multi-organisation workspaces and advanced governance

ENTERPRISE

Contracted

Large organisations, regulated programmes, government agencies and NGOs with contracted requirements.

- Custom quotas agreed per contract
- Dedicated tenancy or self-hosting
- Single sign-on and assurance-grade governance
- Multi-organisation governance for partners, auditors and community contributors
- Premium support with an agreed service level
- Contracted through sales; not available for self-serve checkout

What to know about pricing

Regional pricing

Prices are published in US dollars, Indian rupees, euros and pounds sterling. The India price points are set to familiar local benchmarks rather than a straight currency conversion, and the euro and sterling rates are static approximations reviewed periodically rather than a live feed.

What is not yet live

Annual and multi-year billing periods, nonprofit discounting, usage-based API billing and storage overage billing are stated policy rather than shipped checkout options. Five of the nine published quotas are catalogue entries only and are not yet gated in-product — monitored sites, tracked species and conservation programmes are the three that gate today.

Add-ons are sales-negotiated

Nothing is purchasable as an add-on in-product today. Advanced environmental analytics, premium data sources, enterprise integration suites, sensor and device connectivity, compliance packages, premium support and sector-specific suites are all contracted through sales.

Enterprise, NGO and government shapes

Enterprise, NGO and government tiers are contracted per organisation and are not seeded for self-serve purchase. Professional services — strategy, implementation, data migration, integrations, monitoring enablement, managed operations, governance and assurance, and training — are quoted separately from the platform subscription.

Why teams choose it

Four pillars in one record model, not four tools

Carbon entries, instrument readings, species observations and restoration activities share one workspace, one site reference and one audit trail. Competitors own a single pillar well and leave you to reconcile three to five separate tools; this is the difference that everything else follows from.

Field-to-boardroom evidence lineage

Readings and carbon entries roll into report line items, so an assurance provider can take any number in a published disclosure and trace it back to the record that produced it, rather than to a spreadsheet somebody assembled once.

Native restoration project management

First-class habitat projects with a lifecycle, budget and target hectares, and logged interventions across planting, invasive removal, rewilding, water management, soil remediation, engagement and survey — linked to species observations. No carbon or ESG platform offers this.

Monitoring that reaches the physical world

Nine instrument kinds spanning air, water, weather, soil, energy, flow, acoustic, camera trap and manual sampling, with configurable thresholds evaluated as data lands. Cloud-footprint tools measure what happens in a data centre; this measures what happens in a watershed.

Sensitive-location protection built into the read path

Coordinates for endangered sightings and priority protected sites are generalised on every ordinary read with no way to disable it, and the AI agent surface cannot even form the request for precise values. Protection is structural, not a setting somebody forgets to enable.

An auditable ledger with sourced factors

Emission factors carry their source attribution and valid year, and entries move through draft, recorded and verified, so what has been confirmed is always distinguishable from what has merely been typed in.

Isolation that suits consultancies and NGOs

Every client programme sits in its own workspace with role-scoped access and no cross-tenant writes, so a consultancy runs many isolated programmes from one platform and an NGO collaborates with partners and auditors without widening the blast radius.

Agent-addressable without a bespoke integration

Forty domain tools, four resources and two guided workflows over the same permissions as the interface, so the AI already in your editor or automation stack can do real work on day one rather than after an integration project.

Honest answers to hard questions

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

Because we would rather you knew. There are no paying organisations, no revenue and no case studies, and anyone claiming otherwise about a product in this state is telling you something about how they will report your data later. What does exist is a running production service, a complete web application across all three areas, seeded roles, published legal terms and a public waitlist. Evaluate the product, not a logo wall.

We already have a carbon platform. Why change?

If carbon is genuinely all you need, a specialist may serve you better — several have deeper Scope 3 supplier engagement and AI-assisted factor matching than we do today, and we name that gap on our own comparison pages. The case for changing is if the environmental picture your board asks about is wider than emissions and you are currently reconciling it by hand across three tools.

Your AI section reads more roadmap than product.

Roughly half of it is, and each card says which. Live: the assistant, per-turn cost metering, forty agent tools, threshold evaluation at write time, wildlife-location generalisation and the audit trail. Not built: any AI judgement about your data. We would rather lose a deal on that honesty than win one and have you discover it in week three.

Can we trust an environmental figure that came out of a model?

You should not have to. No AI path can verify a carbon entry or publish a report — both are separate permissioned human actions. The planned drafting work is specified so a line item that cannot name the readings or entries behind it is never written at all. What is still missing is the evaluation set that would measure this, and we say so rather than implying it exists.

Hard questions, continued

We are regulated. Where is your SOC 2?

We do not have one, and we are not going to imply otherwise. SOC 2 Type II and ISO 27001 are planned with no auditor yet engaged. What we can show today is the enforced control set — access control on every operation, tenant isolation, an exportable audit trail, encryption, dependency scanning — and a certification roadmap with dates, owners and honest status per line.

Our data cannot leave our region.

All state currently stays in India by default. That is a real constraint if your requirement is EU or US residency, and we will not pretend a cross-region deployment exists. Dedicated tenancy and self-hosting are available on Enterprise terms, and residency requirements are best raised before a contract rather than after.

Half these quotas are not even enforced.

Correct, and it is published in our own pricing documentation rather than discovered. Monitored sites, tracked species and conservation programmes gate today; instruments, measurements, observations, carbon entries, reports and restoration hectares are catalogue entries awaiting enforcement. In practice that means you are unlikely to be blocked by a limit before launch — but you should know which are real.

Will our field teams actually use it?

Capture is structured and quick — species, method, count, coordinates, time, evidence — and the field role can only add records, never reshape catalogues or retune thresholds, which is what makes it safe to give to a wide crew. The honest gap is that there is no native offline application yet; field capture today happens in the responsive web application.

Where it sits against the alternatives

VS ENTERPRISE CARBON ACCOUNTING PLATFORMS

They measure emissions to an audit-grade standard and carry advisory credibility we do not. We match the ledger discipline — sourced factors, a draft-to-verified lifecycle, scope rollups, eight frameworks — and then keep going into monitoring, biodiversity and restoration that they do not model at all. Where they are clearly ahead: supplier engagement workflows and AI-assisted factor matching, both of which are on our roadmap.

VS SUPPLY-CHAIN AND SCOPE 3 SPECIALISTS

For upstream carbon alone they are the better tool: structured supplier data collection and product carbon footprints are their core and are genuinely deeper than ours. We let you account for Scope 3 in the same ledger as everything else, and we do not have supplier data-request workflows yet. Choose them if the value chain is the whole problem; choose us if it is one part of a wider environmental picture.

VS ENVIRONMENTAL HEALTH, SAFETY AND COMPLIANCE SUITES

They are mature at chemical compliance, risk assessments, incident management and audits — a body of functionality we do not offer and are not close to. We add the carbon ledger, reduction targets, framework disclosure and the whole conservation side on top of a monitoring layer that overlaps theirs. Some organisations will reasonably run both.

VS BIODIVERSITY AND FIELD-SURVEY APPLICATIONS

Community-science apps have vastly larger observer networks and richer species reference data than a young platform can match. What they do not have is enterprise governance, restoration project management, or a route from a sighting to a disclosed figure. We give field capture the same isolation, roles and audit trail as the financial side of environmental work.

VS SUSTAINABILITY MODULES INSIDE A BUSINESS SUITE

If your organisation already runs on that suite, its module inherits the identity, analytics and reporting you have standardised on, which is a real advantage. Our answer is independence: we work with any business system or none, and we measure the impact that happens outside the cloud — sites, instruments, water, air, species and habitat — which those modules do not model.

VS SPREADSHEETS, A MAPPING TOOL AND A SHARED DRIVE

The honest thing about this stack is that it is free, familiar and infinitely flexible, and for a single small programme it can genuinely be enough. It fails at exactly one point: when somebody asks where a number came from and the answer has to survive a reviewer. That is the moment structure, provenance and an audit trail stop being overhead.

START TODAY

Join the waitlist ahead of the 2026 launch

We would rather talk to a few organisations with a real environmental programme than collect a large list. If you have one, tell us about it.

- Join the waitlist to be first in line when onboarding opens
- Read the use-case library and find the two or three that match your programme
- Bring us the awkward question — residency, assurance, a gap on this page — and we will answer it plainly
- Talk to the team about implementation, migration or managed operations if you want help getting there

WEBSITE AND WAITLIST

ecoimpacthub.com

USE CASES

ecoimpacthub.com/use-cases

AI, STATED HONESTLY

ecoimpacthub.com/ai

SERVICES

ecoimpacthub.com/services

TALK TO US

contact@ecoimpacthub.com