

GeoCool-In Digital Strategy: *IT, Web & Geo-DB architecture*

v1.1

(Subscription-Focused)

Version: v1.1 – Board Review Draft

Date: April 2026

Change Log (from v1.0 to 1.1)

- *Made GeoCool-In's core role as a GSHP installer explicit in the introduction.*
- *Clarified that this IT and data strategy evolves GeoCool-In into a data and analytics provider for agricultural professionals.*
- *Highlighted the path toward subscription-based monitoring and analytics services in Executive Summary, Business Context, Value by Phase, and Annex D.*
- *Kept architecture, roadmap, and annex structure stable to ease comparison with v1.0.*

Change Log (from v0.2 to 1.0)

- *Clarified business outcomes and added a board-facing Executive Summary and "Value by Phase" view.*
- *Standardised naming (e.g. earth.geocool-in.com, Compu3-Server) and annex structure (Annex A, B, C, D).*
- *Tightened roadmap timelines and aligned advanced ML work with post-Month 10 evolution.*
- *Moved low-level implementation detail into annexes with explicit references from the main body.*
- *Refined risks, KPIs, and open questions to support board-level decision-making.*

Executive Summary

GeoCool-In is currently a GSHP design and installation company working in agricultural sites (greenhouses, poultry houses, and farms). This IT-Web-DB strategy adds a structured data and analytics layer on top of those installations, so that GeoCool-In becomes not only an installer of efficient systems but also a provider of monitoring, insights, and decision support for agricultural professionals.

Over the next 10 months, the same infrastructure that supports pilots and internal monitoring is deliberately structured as a reusable platform for subscription-based analytics and reporting services (for example environmental dashboards, GSHP performance summaries, and geo-context overlays) that can be sold to farmers, cooperatives, and partners.

This strategy is guided by a Millennial, open-science ethos: data should be understandable, verifiable, and shareable, not locked in black boxes. By making high-quality environmental and GSHP data visible, GeoCool-In contributes to citizen science and broader societal learning around climate, agriculture, and energy.

The roadmap delivers visible value in each phase while building the foundations of this new data and analytics business line:

By Month 2, GeoCool-In will operate a multilingual trust-building one-pager (geocool-in.com) with animated geospatial visuals and basic pilot data.

By Month 4, the first physical monitoring station will be sending data into a production-grade time-series database, with internal dashboards for operations.

By Month 6, the public environmental dashboard (earth.geocool-in.com) will display automated air and weather metrics, demonstrating open, transparent data practices.

Months 7–10 focus on automating data pipelines, offering map exports as a light product, and starting structured ML and geo-event analysis, with deeper GSHP monitoring and ML work continuing beyond this initial horizon.

The architecture follows a clear principle: managed cloud (Railway) for production web and databases, Compu3-Server as an internal analytics and ETL hub, and a common stack (Next.js, Payload CMS, TimescaleDB, QGIS) reused across corporate site, client portals, and public environmental initiatives. This approach minimises early DevOps burden while allowing advanced experiments (ML, scraping, GIS joins) to run safely on internal hardware that underpins future subscription services.

Key decisions requested from the board:

- Confirmation of the 10-month phased roadmap and associated focus (corporate trust first, monitoring and earth. second, advanced ML later).
- Agreement on the public environmental initiative (earth.geocool-in.com) as a strategic asset for reputation, learning, and future analytics products, not a fully monetised product in 2026.
- Support for modest hardware investment (sensors, Raspberry Pis) and for dedicating internal time to Compu3-Server setup, data quality processes, and packaging of monitoring/analytics as optional subscriptions once pilots stabilise.

In summary, this IT and geo-data strategy not only de-risks GeoCool-In's GSHP pilots and improves operational capability, it also creates the technical foundation for a new, subscription-based data and analytics business line in agriculture.

1. Business Context and Goals

GeoCool-In operates at the intersection of GSHP engineering, agriculture, and environmental data. Revenue was originally conceived as primarily coming from project-based GSHP design and installation at individual agricultural sites.

The company needs a digital presence that makes GSHP pilots tangible, make a “technical case” for GSHP, demonstrates data quality, and shows a credible path toward scalable business, now including monitoring and analytics services. The strategic opportunity is to turn those GSHP installations into long-term data nodes and offer ongoing monitoring and analytics subscriptions to agricultural professionals, creating recurring revenue on top of one-off installation projects.

Beyond commercial upside, GeoCool-In aims to contribute to society by treating environmental and performance data as a civic asset: something that farmers, neighbours, researchers, and students can explore and build upon. The open-science and citizen-science approach is part of the brand: transparent data, clear methods, and an invitation to collaborate.

The goals for this 10-month digital strategy are:

- **Trust and visibility:** A high-quality, multilingual corporate site that showcases pilots, team, and partners.
- **Evidence and transparency:** Public environmental dashboards that show real data, methodologies, and open licences.
- **Operational capability:** Internal dashboards and client portals that monitor GSHP performance and site conditions.
- **Future optionality:** A data architecture that supports later products (analytics, reports, ML insights) without locking GeoCool-In into premature complexity.
- **Recurring data services:** Use the shared IT and data architecture to launch optional subscription plans (for example monitoring, reporting, benchmarking) for agricultural clients once the first pilots are stable.

1.1 Commercialisation Path for Data & Analytics

The immediate purpose of this IT strategy is to support GeoCool-In’s GSHP pilot projects and credibility. In parallel, it deliberately builds the components needed for a future subscription-based business line:

- Standardised data ingestion (sensors plus public APIs) into TimescaleDB.
- Reusable dashboards and map components that can be parameterised per client or site.
- An internal analytics hub (Compu3-Server) where new analytics and ML features can be prototyped and then rolled into repeatable offerings.

From 2027 onward, this enables tiered services for agricultural professionals, ranging from basic monitoring subscriptions to higher-value analytics and geo-context products, all running on the same core platform.

2. Value by Phase (*Board View*)

The roadmap is structured so that each phase delivers concrete, externally visible value while building the underlying platform for future subscription services.

Phase	Timeframe	External value (board / partners)	Internal capability unlocked
0	Month 0	Clear visual identity and story for GeoCool-In	Brand, Figma, QGIS SVG assets ready
1	Months 1–2	Live multilingual one-pager with animated map and contact forms	Next.js/Payload stack on Railway, basic data model
2	Months 3–4	First live pilot sensor connected (proof of monitoring ability)	Compu3-Server foundations, sensor schema and ingestion
3	Months 5–6	Public earth.geocool-in.com showing automated air/weather data and service-like UX	Reusable dashboard components, cron-based API ingestion
4	Months 7–8	"Export your map" feature and richer public maps; potential lightweight paid map exports	Automated ETL, backup flows, map export pipeline
5	Months 9–10	Early ML-backed geo/GSHP insights for storytelling and premium analytics pilots	Structured ML experiments on Compu3-Server and beyond

Beyond Month 10, the same architecture supports scaling GSHP monitoring to more sites, adding supervisory control logic, and deepening geo-event and ML analysis.

These longer-term trajectories are outlined in Annex C (earth. EcoWatch) and Annex D (GSHP Monitoring System), and can be progressively packaged into subscription offerings.

3. Objective

Create a coherent digital architecture where GeoCool-In's field data, public environmental datasets, and geospatial layers flow into:

- A trust-building corporate site (geocool-in.com) to support pilots, fundraising, and partnerships.
- An open environmental dashboard (earth.geocool-in.com) that showcases transparent, high-quality environmental data.
- Internal operations and client monitoring surfaces (inside.geocool-in.com and clients.geocool-in.com) for GSHP performance and site management.

Over 10 months, sequence quick visual wins (QGIS SVG maps, simple dashboards) and core infrastructure (TimescaleDB on Railway, Compu3-Server as analytics hub), then progressively layer in automated ETL, public API ingestion, and targeted ML experiments, so that GeoCool-In can evolve from purely project-based GSHP installation towards a recurring data and analytics business for agricultural professionals.

4. Scope

4.1 In Scope (2026)

- Corporate multilingual one-pager at geocool-in.com, including animated SVG maps and basic pilot site storytelling.
- Internal dashboards and later client portal (inside.geocool-in.com, clients.geocool-in.com) for GSHP monitoring and site metrics.
- Public environmental dashboard at earth.geocool-in.com, initially focused on air quality (PM2.5), temperature (air and soil), and water pollutants.
- Hardware procurement and deployment of monitoring sensors (ESP32-based pods, Raspberry Pis, DS18B20, SHT31/DHT22, PM2.5 sensors, Pi camera) at pilot sites.
- Data layer based on PostgreSQL/TimescaleDB on Railway (production) with an analytics replica on Compu3-Server.
- QGIS projects (geology, soil, land use, cadastral/property) and QGIS→SVG export pipeline for web visuals.
- CSV-first manual data uploads followed by scripted scraping and open API ingestion.
- Research-grade ML and analytics experiments on Compu3-Server (and companion machines) focused on time-series modelling, geo-events, and GSHP analysis.

4.2 Out of Scope (for this v1 horizon)

- Production-grade, fully automated MLOps for all models.
- Full self-hosting of the production web stack and databases (Railway remains primary production host in 2026).
- Non-GSHP digital products unrelated to environmental monitoring and geothermal applications.

5. Current State

5.1 Web and Database Stack

- Proposed stack: Next.js App Router + Payload CMS + TimescaleDB deployed on Railway, with Grafana and Umami as supporting services.
- Planned apps share a single monorepo and database: geocool-in.com (corporate site), inside.geocool-in.com (internal dashboards), clients.geocool-in.com (client portal), and earth.geocool-in.com (public environmental dashboard).

5.2 Monitoring Station and Hardware

- Sensors (ESP32, Raspberry Pis, DS18B20, SHT31, PM2.5 sensor, Pi camera) are specified but not yet deployed at pilot sites.
- Initial deployments target greenhouses and a chicken house, where environmental variability and harsh conditions will stress-test the monitoring approach.

5.3 Internal Infrastructure

- Compu3-Server (HP i7-8550U, Ubuntu Server LTS, PostgreSQL, SSD plus HDD) is available but largely idle.
- A high-performance personal workstation (Compu1) with GPU/NPU capabilities is also available for heavier ML experiments.

5.4 Design and GIS Assets

- Branding and Figma wireframes for the one-pager are in progress and expected mid-April.
- QGIS layers for geology, soil, and land use are prepared and ready for SVG export via a QGIS Python pipeline.
- Daily weather data sources (for example CWA, EPA) are identified for future ingestion.

6. Problems and Constraints

- **Design bottleneck:** Web development is blocked until Figma and branding are finalised; Elsa cannot implement what is not yet designed.
- **Cost and simplicity:** The public-facing site must remain low-complexity and low-cost in 2026.
- **Sensor deployment lag:** Hardware procurement and on-site installation will take time, so early visualisation must rely on historical CSVs, sample data, and QGIS static exports.
- **Operational risk:** Compu3-Server must create tangible value without becoming an operational burden or single point of failure for production traffic.
- **Ambition vs capacity:** Advanced ML on environmental and GSHP data is strategically important but time-limited; it must be staged behind fundamentals (monitoring, ETL, dashboards).

7. Target State

7.1 High-Level Architecture

- **Railway (Production):**
 - Hosts the monolithic Next.js application serving geocool-in.com, clients.geocool-in.com, inside.geocool-in.com, and earth.geocool-in.com.
 - Runs the master PostgreSQL/TimescaleDB instance for all production data.
- **Compu3-Server (Internal Analytics Hub):**
 - Runs scheduled Python jobs (scraping, ETL, backups) and can host JupyterLab sessions.
 - Maintains an analytics replica of production data for heavy queries, GIS joins, and ML experiments.
 - Stays off the public internet; no production traffic is routed through Compu3-Server.

7.2 Data Flow Overview

- **Sensor data:** On-site ESP32 pods and Raspberry Pis send sensor readings to the ingestion endpoints on Railway (TimescaleDB), using a consistent schema for environmental and GSHP metrics.
- **Public APIs:** Compu3-Server runs cron jobs that fetch Taiwan EPA, CWA, and other open data, validate them, and push cleaned data into dedicated schemas on Railway.
- **Applications:** geocool-in.com, clients.geocool-in.com, inside.geocool-in.com, and earth.geocool-in.com all query Railway for their data needs, reusing shared UI components for charts and maps.

8. Strategic Principles

1. **Managed cloud for production, Compu3-Server for analytics and experimentation.** SLA-critical workloads run on Railway, while Compu3-Server handles scraping, ETL, heavy GIS joins, and ML in a controlled environment.
2. **CSV to API to ML pipeline.** Start with manual CSV uploads and simple scripts to prove value; then progress to automated API scrapers and eventually ML inference, avoiding premature complexity.
3. **Visuals first.** Use QGIS→SVG exports to deliver compelling map-based visuals quickly on the corporate site and earth.geocool-in.com, even before full sensor networks are live.
4. **AI as augmentation, not magic.** Focus on anomaly detection, summarisation, and classification that directly support operations and storytelling, rather than speculative large-scale models.
5. **Reuse and standardisation.** Reuse Next.js components, charting libraries, and schemas across corporate, client, and public surfaces to minimise maintenance and ensure consistency.
6. **Millennial, open-science ethos.** Default to open-software, open-data where possible: explain methods, cite data sources, and design earth.geocool-in.com so that citizens, students, and researchers can reuse and question the data, not just passively consume charts. Aim at becoming the air-quality reference app (market niche exist).

9. Proposed Architecture and Operating Model

9.1 Web Surfaces

- **Corporate site (geocool-in.com):**
 - o Multilingual one-pager with pilot stories, animated SVG maps, and concise explanation of GSHP and environmental monitoring.
 - o Prioritises load performance, clarity, and credibility over complex interactivity.
- **Public environmental dashboard (earth.geocool-in.com):**
 - o Dedicated to air quality, temperature (soil and air), and water quality metrics, using the same chart components as internal dashboards.
 - o Designed as an open-science instrument: dark theme, clear data sources, timestamped readings, and open licences, aligned with a Millennial ethos of transparency and citizen science participation.
 - o Initially unlisted or softly promoted, but shareable with partners and community, and ready to be packaged as part of monitoring subscriptions.

- **Internal and client portals (inside.geocool-in.com, clients.geocool-in.com):**
 - o Real-time GSHP monitoring, alarms, and historical views for internal team and private clients.
 - o Built later in the roadmap as stable ingestion from pilot sites is confirmed, with a view to offering them as paid monitoring and analytics portals.

9.2 Data and Analytics Layer

- **TimescaleDB on Railway:** master time-series store for sensor and environmental data.
- **TimescaleDB on Compu3-Server:** analytics replica for heavy aggregations and GIS joins against QGIS-exported data.
- **JupyterLab and Python stack on Compu3-Server:** for exploratory analysis, model prototyping, and ETL development, supporting the creation of value-added analytics features for future products.

10. Roadmap by Phase (10-Month Plan)

Phases are defined as: Phase 0 (Month 0, prerequisites), Phase 1 (Months 1–2), Phase 2 (Months 3–4), Phase 3 (Months 5–6), Phase 4 (Months 7–8), Phase 5 (Months 9–10).

Each phase delivers at least one visible outcome to investors and partners while building long-term capability for subscription-based services.

Phase 0 – Immediate Prerequisites (Month 0)

- Finalise logo, brand colours, and core copy for the one-pager.
- Run QGIS→SVG export script to generate geology, soil, land use, and pilot site map layers.
- Complete Figma wireframes for the one-pager, placing the SVG map as a central visual element.
- Order ESP32, DS18B20, Raspberry Pis, PM2.5 sensors, SHT31, and one Pi camera.

Phase 1 – Corporate Trust & Compu3-Server Foundations (Months 1–2)

- Build and deploy the Next.js/Payload one-pager (four languages) on Railway.
- Implement animated SVG map in the hero section as a quick visual win.
- Set up basic access to Compu3-Server (SSH plus tmux) and install local PostgreSQL/TimescaleDB (see Annex A).
- Manually upload historical environmental and water CSVs via Payload to test data structures and dashboards.

Phase 2 – Sensor Ingestion & Dashboarding (Months 3–4)

- Deploy the first Raspberry Pi pilot site (for example “Templo live recording sessions”).
- Implement TimescaleDB sensor schemas on Railway and Next.js ingestion routes.
- Launch inside.geocool-in.com with Grafana for internal GSHP and environmental monitoring.
- Set up JupyterLab on Compu3-Server and write initial Python scripts for parsing CSVs and prototyping EPA/CWA scraping.

Phase 3 – Client Portal & earth. Launch (Months 5–6)

- Build clients.geocool-in.com using Tremor.js/Recharts for investor and client views.
- Reuse the same UI components to launch earth.geocool-in.com as a public view for temperature, relative humidity, PM2.5, and basic water metrics.
- Go live with earth.geocool-in.com using a mix of live Pi sensor data and manual CSV uploads as needed, creating the first end-to-end experience that resembles a subscription monitoring service.
- Deploy cron jobs on Compu3-Server that fetch Taiwan EPA/CWA APIs and push cleaned data into Railway on a daily schedule.

Phase 4 – Automation & Map Exports (Months 7–8)

- Build the "Export your map" feature on earth.geocool-in.com, allowing users to select a bounding box and receive a printable PNG/PDF map layered with GeoCool-In branding.
- Treat map export as an experiment for lightweight paid services (for example branded environmental/geo reports for farms or partners).
- Fully automate ETL pipelines: Compu3-Server runs backup and ingestion flows, pulling daily snapshots from Railway and pushing processed aggregates back.

Phase 5 – Advanced Geo & ML Foundations (Months 9–10)

- Begin structured NLP experiments on Compu3-Server (and optionally Compu1) to scrape and classify news related to conflict, mining, and environmental events.
- Update earth.geocool-in.com to support ML-derived heatmaps (for example mining concessions, geo-events) via GeoJSON APIs.
- Run time-series forecasting on soil temperatures versus ambient weather and GSHP operation to identify patterns and anomalies that can feed into premium analytics offerings.
- Prepare results and visualisations as narrative prototypes, not yet full products; more advanced ML and GSHP monitoring work continues beyond Month 10 (see Annex C and D).

11. Risks and Dependencies

- **Figma and branding bottleneck**
 - o Risk: Elsa's work is blocked until Figma designs are stable.
 - o Mitigation: Timebox design/branding to Phase 0; lock v1 scope and treat later design changes as Phase 3+.
- **Hardware delays and deployment friction**
 - o Risk: Procurement, networking, and environmental conditions may delay sensor commissioning.
 - o Mitigation: Use historical CSVs and mocked data to drive early dashboards; keep a minimal viable sensor set for first pilots.
- **Scope creep on earth.geocool-in.com**
 - o Risk: Overloading the public dashboard with features before core GSHP monitoring is solid.
 - o Mitigation: Keep earth. focus on a small set of metrics and basic map exports during 2026; advanced geo/ML features remain experimental.

- **Compu3-Server operational overhead**
 - o Risk: Compu3-Server could become hard to maintain.
 - o Mitigation: Restrict exposure (no public services), automate basic maintenance, and document minimal operating procedures.

12. KPIs and Success Criteria

12.1 External / Board-Level KPIs

- By Month 2: Live multilingual corporate site with functional contact form and animated map.
- By Month 4: At least one pilot site sending live sensor data to TimescaleDB without repeated packet loss.
- By Month 6: earth.geocool-in.com displaying automated daily PM2.5 and temperature data, without manual intervention.
- By Month 10: First ML-backed insights (for example anomalies, basic geo-event maps) demonstrated as part of investor/partner storytelling and as prototypes for future paid analytics.

12.2 Internal Technical KPIs

- Data uptime and quality for each sensor pod (for example percentage of readings passing QC checks).
- Time to commission a new site (from sensor shipment to live dashboards).
- ETL automation coverage (percentage of critical flows running via cron/automation without manual steps).
- Number of stable, re-usable dashboard components and map layers used across web surfaces.

13. Open Questions

- Map export implementation: For the "Export your map" feature, should a local tile server be hosted on Railway, or will composition stay purely SVG/browser-based in 2026?
- Prioritisation of ML tracks: Which ML experiments (time series, geo-events, raster analysis) are most aligned with 2027 funding and product strategy, and should therefore be prioritised in Months 9–10?
- Subscription design: What tiers and price points (basic monitoring vs advanced analytics) make sense for agricultural clients once the platform is stable?

Annex A – Compu3-Server Internal Configuration (Summary)

Compu3-Server is the internal analytics worker node.

It pulls data from Railway, runs ETL and ML jobs, and pushes processed metrics back, without ever serving public web traffic.

Key points (details retained from v0.2):

- Ubuntu Server LTS, SSH key-based access only, with firewall restricting inbound connections to SSH from trusted networks.
- Use of tmux for long-running scraping and ETL scripts to avoid reliance on a single SSH session.
- JupyterLab environment on Compu3-Server for local data analysis, accessible from development machines.
- Local PostgreSQL/TimescaleDB instance acting as analytics replica, updated via scheduled pg_dump/restore or similar mechanisms.

Annex B – QGIS → SVG Export Pipeline for Web Maps

This annex describes the bridge between QGIS spatial data (soil, geology, vegetation, land use) and web-ready animated map visuals.

[^1]

Key elements:

- A Python script inside QGIS exports each layer as transparent SVG at a standard canvas size and extent, generating a manifest file used by the web map component.
- The web component stacks SVG layers (base, geology, ecology, operations) with CSS animations and optional live sensor overlays.
- Project CRS, styling, and export order are standardised to ensure consistent visuals across pages.

Annex C – earth.geocool-in.com "EcoWatch" Concept

earth.geocool-in.com ("EcoWatch") is the public environmental monitoring page and internal learning lab, explicitly grounded in a Millennial open-science and citizen-science ethos. It is inspired by open-science dashboards and notebook-style explorations.

Highlights:

- Combines external public feeds (EPA, CWA, MOEA water quality) with GeoCool-In's own sensors.
- Presents a hero KPI strip, map section, timeseries panel, water quality panel, photo gallery, and open data footer.
- Requires incremental additions to schema and storage (for example external readings table, photo storage via Cloudflare R2 or MinIO on Compu3-Server).
- Serves as a bridge between rigorous internal monitoring and transparent public engagement.

Longer-term ML work on news, maps, and images (geo-event detection, concession mapping, influence maps) builds on this foundation and extends beyond the 10-month roadmap, and can feed future premium analytics tiers.

Annex D – GSHP Monitoring System – Strategic Deployment Plan (2026–2027)

The GSHP Monitoring System annex outlines a multi-year plan for research-grade monitoring across agricultural sites.

In addition to supporting internal R&D and project performance verification, this monitoring platform is designed to be productised as a standard monitoring and analytics service that can be offered on a subscription basis to agricultural clients.

Key points:

- Phased rollout of environmental pods, HP manifold nodes, and soil monitoring nodes, starting with a small number of reference sites.
- Standardised hardware, MQTT topic structures, and JSON schemas to support repeatable deployments.
- Strategic principles that emphasise calibration, data quality, and metadata over raw sensor count.
- KPIs focused on data uptime, commissioning time, COP/SPF uncertainty, and the number of sites and supervisory rules in production.
- Architecture (edge pods plus site gateway plus central platform) intentionally designed so it can be replicated across sites and packaged as a recurring monitoring and analytics service, not just a one-off pilot tool.