GenAI for Geospatial Challenge

EMEA Edition

End of Challenge Report



14	48	$650k	3
winning teams	proposals	AWS cloud credits	regions

aws thriveGEO WOMEN+ IN GEO Lunate AI

Table of Contents

Foreword 4

Executive Summary 5

The GenAI for Geospatial Challenge - EMEA Edition 6

Challenge Overview 7

- The Launch Phase: Oct-Dec 2025 7
- The Award Phase: Dec-Jan 2026 8
- The Build Phase: Jan-May 2026 10
- The Show Phase: Jun-Jul 2026 11

Meet the GenAI for Geospatial Challenge Winners 14

- Aspia Space 16
- GEOTERRAIMAGE 18
- NatureAlpha 20
- Ocean OS 22
- Metaplanet 24
- VIDA 26
- VISTA 28
- Urban Climate MENA 30
- ATOM AI 32
- BigData Ghana 34
- RASID 36
- Sproxil 38
- BolgiaTen 40
- Globeholder AI 42

Key Takeaways 44

What's Next 47

Acknowledgements 48

This report shares the highlights, insights and impact of the GenAI for Geospatial Challenge – EMEA Edition, 2025-2026.

Foreword

"GeoAI is here. The innovators proved it. The GenAI Geospatial Challenge gave startups and scaling businesses the platform, the AI tools, and the community to show that the future of geospatial intelligence – from defence to climate to finance – belongs to those who move fastest. The condition of GeoAI in EMEA is not emerging. It is thriving."

Generative AI is opening new possibilities for how people access, apply and make decisions based on geospatial information. But technology alone is not the story. The real opportunity lies in how innovators combine new AI capabilities with domain expertise, responsible data practices, cloud infrastructure and a clear understanding of real-world user needs.

Building on the strength of the challenge's previous edition, launched in North America in 2024, AWS and thriveGEO - together with partners Women+ in Geospatial and LunateAI - brought the Generative AI for Geospatial Challenge to Europe, the Middle East and Africa for the first time. This EMEA edition brought together innovators, industry leaders and cloud resources from across the region to explore and expand the art of the possible for big-impact applications using geospatial data.

Over the course of the challenge, we saw 14 teams move beyond abstract ideas and begin developing practical solutions for climate resilience, environmental monitoring, urban systems, infrastructure planning, decision intelligence and data access. These projects do not suggest that AI can solve complex environmental or societal challenges on its own. Rather, they show where AI can act as an enabling layer: helping people ask better questions, work with complex datasets, prototype faster and translate technical outputs into forms that more people can understand and act on.

This report captures the outcomes of the GenAI for Geospatial Challenge – EMEA edition. It sets out the teams, themes and early prototypes that emerged, along with the lessons learned about what it takes to move from promising ideas to credible, usable solutions.

The challenge marks only the beginning. Its real value lies in what it revealed: a growing community of innovators across the region ready to experiment with generative AI responsibly, creatively and with a genuine connection to solving real societal challenges.

We hope this report offers a useful window into where GenAI for geospatial is heading, and an invitation not only to observe the future of geospatial AI, but to help build it.



Julia Wagemann, PhD
Co-CEO, thriveGEO



Phil Cooper
Geospatial Global Tech BD, AWS

Executive Summary

The GenAI for Geospatial Challenge – EMEA Edition set out to answer a practical question: what happens when bold ideas, generative AI and geospatial data are given the time, cloud resources and expert support to become real? Over eight months, from launch to showcase, **14 winning teams** from across Europe, the Middle East and Africa answered that question with working prototypes addressing some of the region's most pressing challenges – from climate resilience and urban planning to ocean health, forest traceability and public health.

Building on the strength of the challenge's first edition, launched in North America in 2024, AWS and thriveGEO brought the initiative to EMEA for the first time, with support from Women+ in Geospatial, LunateAI and Cloudscaler. The challenge opened for proposals in October 2025 at a launch event in London, drawing **48 submissions** from across the region. A judging committee spanning expertise in Generative AI, geospatial technology, cloud computing and venture building selected 14 teams to take part in a **five-month Build Phase**, supported by **$650,000 worth of AWS cloud credits**, technical guidance and a structured programme of community support.

Beyond the technical diversity of the solutions, the cohort reflected the geographic and cultural breadth of the EMEA region, bringing together innovators from Europe, the Middle East and Africa. The winning teams also included strong female representation, with women participating in 11 of the 14 teams, including one all-woman team.

The result is a portfolio of prototypes spanning four themes: **Climate Resilience & Environmental Intelligence; Infrastructure & Smart Cities; Data Foundations, Trust & Scalable Platforms; and Ecosystems, Markets & Decision Intelligence**. Teams built agentic assistants that turn plain-language questions into spatial analysis, foundation models trained on satellite and ocean data, and copilots that help insurers, regulators and planners act on geospatial evidence with greater speed and confidence. By June 2026, these prototypes were being presented to the wider industry at the AWS Geospatial Day in London and at a final Demo Day, with outcomes reaching even wider audiences through a dedicated podcast series, as well as available for the wider AWS teams to see and discuss with their customers and partners.

Taken together, the cohort's progress points to a clear conclusion: **generative AI works best not as a replacement for geospatial expertise, but as an interface to it** – helping people ask better questions, navigate complex datasets and turn technical outputs into decisions they can trust. That promise depends on strong data foundations, the right cloud infrastructure, and a continued focus on making powerful tools genuinely usable by the people who need them.

This report tells the story of that journey in full: how the challenge was launched, judged and implemented; what each of the 14 teams built and learned; the patterns now shaping geospatial AI across the region; and what it will take to carry these prototypes from promising demonstrations to tools in real-world use.

The GenAI for Geospatial Challenge

EMEA Edition

14 winning teams

48 proposals

$650k AWS cloud credits

3 regions

About the challenge

The GenAI for Geospatial Challenge is an innovation initiative that invites innovators and practitioners to submit bold ideas for big-impact geospatial applications using generative AI. It was first launched in 2024 by the Taylor Geospatial Institute (TGI) and AWS, with a regional focus on North America. Building on the strength of that first edition, AWS and thriveGEO brought the challenge to Europe, the Middle East and Africa, with the aim to inspire and support the development, commercialisation and scaling of innovative geospatial applications across the region.

To support these ideas, selected teams received a total of $650,000 worth of AWS cloud credits spread across projects, AWS landing-zone support for enterprise-ready scaling and technical support from AWS experts. Selected teams were also provided with community support, visibility and recognition.

The EMEA edition was hosted by thriveGEO and AWS, with support from Women+ in Geospatial, LunateAI and Cloudscaler.

01

Launch
Oct-Dec 2025

Challenge was launched at the Royal Geographical Society in London.

02

Award
Dec 2025

Evaluation of submitted proposals to select the 14 winning teams.

03

Build
Jan-May 2026

Selected teams transformed their ideas into working prototypes using AWS credits and support.

04

Show
Jun-Jul 2026

Outcomes are celebrated and presented.

Challenge Overview

The Launch Phase

Royal Geographical Society
London - 8 October 2025

100+ attendees — innovators, cloud architects, practitioners and investors

Oct-Dec 2025

The challenge officially opened for proposal submissions on 8 October 2025, at a launch event held at the Royal Geographical Society in London. The event brought together more than 100 geospatial innovators, cloud architects, practitioners and investors from across Europe, the Middle East and Africa, marking the start of the submission window that ran through to early December.

Alongside the launch, the day offered participants a first look at how generative AI and cloud technology could be applied to geospatial challenges, with contributions from AWS and guest speakers setting the scene for what the challenge hoped to achieve. The event brought together speakers and panellists from across the geospatial, cloud, investment and innovation ecosystem. A full list of contributors is included in the Acknowledgements section on page 48.

The event also linked the new EMEA edition to its North American predecessor, with Sparkgeo and LuxCarta sharing how their Challenge projects evolved beyond the Build phase.

> **"Geospatial AI programmes need to reach beyond usual networks, inviting more regions, backgrounds and lived experiences to shape the questions and tools. Through Women+ in Geospatial, we helped bring more women+ and EMEA perspectives into the challenge."**
>
> – Mariam Crichton, Strategic Partnerships Director, Women+ in Geospatial

Fig. 01



Fig. 01-02-03

A few moments from the AWS launch in London on 8 October 2025.

Fig. 02

Fig. 03



The Award Phase

Dec 2025

Following the close of submissions, proposals were reviewed by a judging committee with expertise across generative AI, geospatial technology, cloud computing, foundation models, Earth observation, scientific innovation, product development and venture building.

48 → 14

From 48 proposals, 14 teams were selected to take part in the Build Phase.

The judging committee



Phil Cooper
Geospatial Global Tech BD, AWS



Nadine Alameh, PhD
Founder, LunateAI



Bruno Sanchez-Andrade Nuño, PhD
Chief Science Officer and Co-Founder, LGND AI; Co-Founder, Clay



Charlette N'Guessan
Data Solutions and Ecosystem Lead, AMINI; Founder, Africa-Centred Scientific Publication "Afrokwary"



Julia Wagemann, PhD
Co-Founder, thriveGEO; Founder and Director, Women+ in Geospatial



Mounia El Baz
EO Digital Innovation Engineer, European Space Agency

How proposals were assessed

Each proposal was judged against six criteria, covering both the ambition and the practical viability of the idea:

01 — **Potential impact**
Will the idea improve the human condition or society?

02 — **Technical sophistication**
Does the solution show a strong understanding of state-of-the-art GenAI?

03 — **Commercial implications**
Is there a clear market and pathway to commercialisation?

04 — **Scalability**
Can the idea scale beyond a one-off prototype?

05 — **Integration of geospatial data**
Is location intelligence essential to the solution?

06 — **Team strength**
Does the team have the expertise to execute?

What the winners received

AWS cloud credits
Prototype development and scaling

Landing-zone support
from AWS partner Cloudscaler

AWS technical support
Technical guidance and Q&A

Community & visibility
thriveGEO, AWS and partner channels

Storytelling opportunities
Spotlights, podcasts and final report

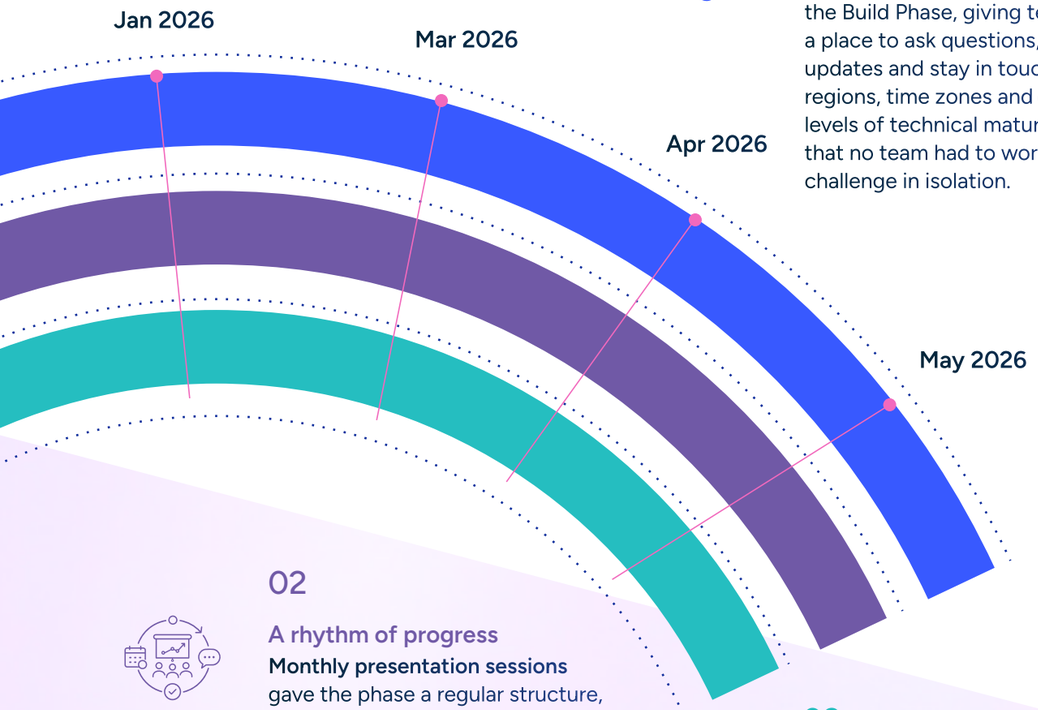
The Build Phase

Jan-May 2026

The Build Phase gave the 14 winning teams five months to turn their proposals into working prototypes, supported by AWS cloud credits, technical guidance and a structured programme of community support.

The Build Phase support system

Guidance, check-ins and peer support helped teams move from early concepts to functioning prototypes.



01

Staying connected

A private Slack workspace kept the cohort connected throughout the Build Phase, giving teams a place to ask questions, share updates and stay in touch across regions, time zones and different levels of technical maturity - so that no team had to work through a challenge in isolation.

02

A rhythm of progress

Monthly presentation sessions gave the phase a regular structure, with winners sharing progress, surfacing blockers and discussing practical questions on data, architecture, cost and user validation. These sessions also created a direct route to raise technical questions on credits, cloud setup and implementation with the AWS team.

03

Building expertise together

As agentic systems emerged as a recurring theme across the cohort's prototypes, AWS led a **dedicated deep-dive session** on the topic, helping teams understand how agentic AI is being applied in practice today and where it can create genuine value in geospatial workflows.

> "thriveGEO worked closely with the challenge teams to turn their winning ideas into functional prototypes. Together with AWS, we connected them with the resources and expertise they need to succeed. We are incredibly proud of what the teams achieved with the AWS credits which they have been awarded."
>
> – Sabrina Szeto, thriveGEO

The Show Phase

From early concepts to functioning prototypes

The Show Phase gave teams space to present their work, share lessons and connect prototypes to real users.

Jun-Jul 2026

The Show Phase brought the challenge to a close, giving teams the opportunity to present their prototypes to the wider geospatial community and share what they had learned along the way.

GenAI for Geospatial Challenge Panel, AWS Geospatial Day, London

The AWS Geospatial Day on 2 June 2026 provided the first public platform to share outcomes from the challenge with a wider geospatial and cloud innovation audience. Selected teams presented their prototypes, explained the problems they were addressing and showed how generative AI was being applied in practice. Discussions extended beyond the technical outputs to questions of maturity, user needs, data access, scalability and the role of cloud infrastructure in moving from early concept to operational product.

Fig. 04



Fig. 05



Fig. 04-05

A few moments from the AWS Geospatial Day in London on 2 June 2026, where selected GenAI for Geospatial Challenge teams shared their prototypes and joined a panel discussion on moving from early ideas to working geospatial AI solutions.

Demo Day

At the end of the challenge, teams came together for a final Demo Day, presenting what they had built and reflecting on the journey from proposal to prototype. The showcase highlighted the breadth of progress across the cohort, including operational services, new models and strong early interest from potential customers.

Fig. 06

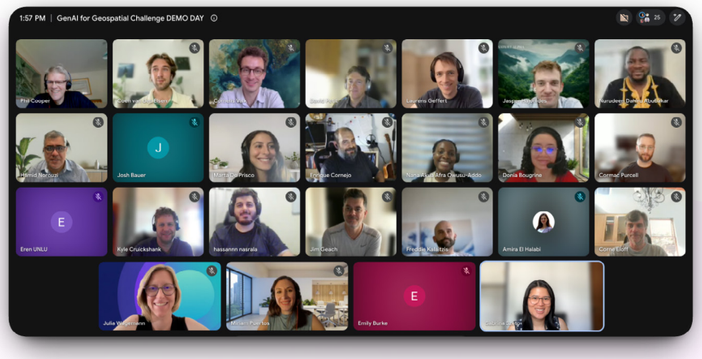


Fig. 06

Final Demo Day for the GenAI for Geospatial Challenge, bringing the cohort together online to present their prototypes, share progress and reflect on the journey from proposal to working solution.

Scene From Above podcast

To extend the reach of the challenge beyond live events, its outcomes are also being shared through a dedicated season of the Scene From Above podcast. The series spotlights winning teams from across the EMEA cohort, giving listeners a closer look at the ideas, prototypes and regional contexts behind the challenge. Each episode explores how teams applied generative AI to geospatial problems, what they learned during the Build Phase and where they see opportunities for further development.

Episode 1

Africa Team Spotlight with **Nana Akua Afra Owusu-Addo, BigData Ghana**

Episode 2

Middle East Team Spotlight with **Reda Haidar, RASID**

Episode 3

Europe Team Spotlight with **Cormac Purcell, BolgiaTen**

Episode 4

Report Launch with **Phil Cooper, AWS, and Nadine Alameh, LunateAI**

> "This is the real power of AI: empowering innovators anywhere to transform complex data into scalable actionable solutions for our world!"
>
> – Nadine Alameh, Founder, LunateAI

Meet the GenAI for Geospatial Challenge Winners

The 14 winning teams reflect the breadth of the challenge, spanning climate resilience, urban systems, data infrastructure and decision intelligence. Their projects naturally group into four themes:

Climate Resilience & Environmental Intelligence



As climate risks accelerate, static monitoring is no longer enough. These projects use GenAI and Earth observation to detect patterns, model uncertainty and turn environmental data into actionable intelligence.

PROJECTS: Aspia Space • GEOTERRAIMAGE • NatureAlpha • Ocean OS

Infrastructure & Smart Cities



Cities and infrastructure systems are becoming more complex and exposed to climate, mobility and operational risk. These teams explore dynamic geospatial intelligence for planning, routing, risk and scenario analysis.

PROJECTS: MetaPlanet • VIDA • VISTA • Urban Climate MENA

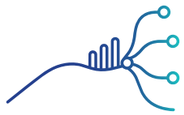
Data Foundations, Trust & Scalable Platforms



Behind every applied AI system is an invisible data layer. These projects focus on making geospatial data reliable, interoperable, trustworthy and ready for AI at scale.

PROJECTS: ATOM AI • BigData Ghana • RASID • Sproxil

Ecosystems, Markets & Decision Intelligence

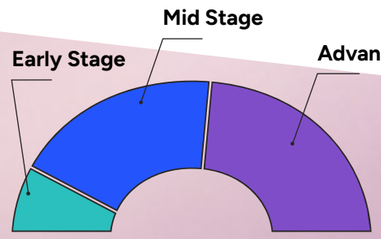


AI becomes most useful when it supports real-world decisions. These projects connect geospatial intelligence with markets, policy, trade and risk-based decision-making.

PROJECTS: BolgiaTen • Globeholder AI

Winning Teams at a Glance

	TEAM	LEAD PROPOSER	PROJECT SYNOPSIS
01	Aspia Space	James Geach	Probabilistic EO scenario planning for climate resilience.
02	GEOTERRA	Kyle Cruickshank	Chat-based climate insights for rainfall, agriculture and livelihoods.
03	NatureAlpha	Jasper Hajonides	Foundation models for ecosystem and nature risk.
04	Ocean OS	Laurens Geffert	Ocean intelligence for offshore wind and environmental monitoring.
05	MetaPlanet	David Petit	AI-native site risk intelligence for insurance decisions.
06	VIDA	Philippe Raisin	Conversational geospatial agent for infrastructure and planning.
07	VISTA	Cornelis Valk	Agentic geospatial AI for evidence-based spatial decisions.
08	Urban Climate MENA	Hamid Norouzi	Urban heat intelligence for city-scale climate and health planning.
09	ATOM AI	Josh Muranga Bauer	AI-ready geospatial datasets and validation workflows.
10	BigData Ghana	Nana Akua Afra Owusu-Addo	Forest loss detection and cocoa supply-chain transparency.
11	RASID	Reda Haidar	Agentic EO workflows for methane detection and monitoring.
12	Sproxil	Nurudeen Abubakar	Settlement-based malaria targeting and health planning.
13	BolgiaTen	Cormac Purcell	Maritime intelligence for dark-vessel, spoofing and jamming risk.
14	Globeholder AI	Eren Unlu	Urban-region embeddings and natural-language geospatial intelligence.



Concept refined • Prototype in progress

Working prototype • Demo-ready

Pilot-ready • Being tested with users or partners • Already integrated into an existing product or service

Case Study

Aspia Space

ManyWorlds

Trustworthy, data-driven probabilistic scenario planning to improve climate resilience.

CHALLENGE

Soil moisture is a critical variable for agriculture, insurance, carbon and climate modelling, yet existing monitoring is often sparse, expensive or too local to support field-scale decisions.

APPROACH

Aspia Space used generative Earth observation models, Sentinel-1 SAR, synthetic cloud-free Sentinel-2, meteorological data and ground-truth probes to estimate UK soil moisture, with uncertainty via conformal prediction.

WHAT WAS BUILT

The team built a working ManyWorlds prototype for soil moisture prediction. The model produces continuous volumetric water content time series at 5-30 cm soil depth and 10 m resolution, fine enough to show variation across individual fields.

RESULTS

The prototype is already generating pilot data across representative UK locations. It is now being scaled to cover the UK mainland, with lightweight dashboard tools being developed for agronomists and other users.

AWS IMPLEMENTATION

AWS Batch was used to run large modelling jobs in parallel. Heavy workloads used c7i.8xlarge instances: CPU-optimised AWS machines with 32 virtual CPUs and 64 GiB RAM, suited to high-resolution weather and soil-moisture modelling.

NEXT STEPS

Continue technical development, improve the demo, seek pilot partners, integrate the capability into an existing product and explore commercial opportunities with agronomy, insurance and carbon-market users.

REGION: Europe • United Kingdom

STAGE: Advanced • Pilot-ready

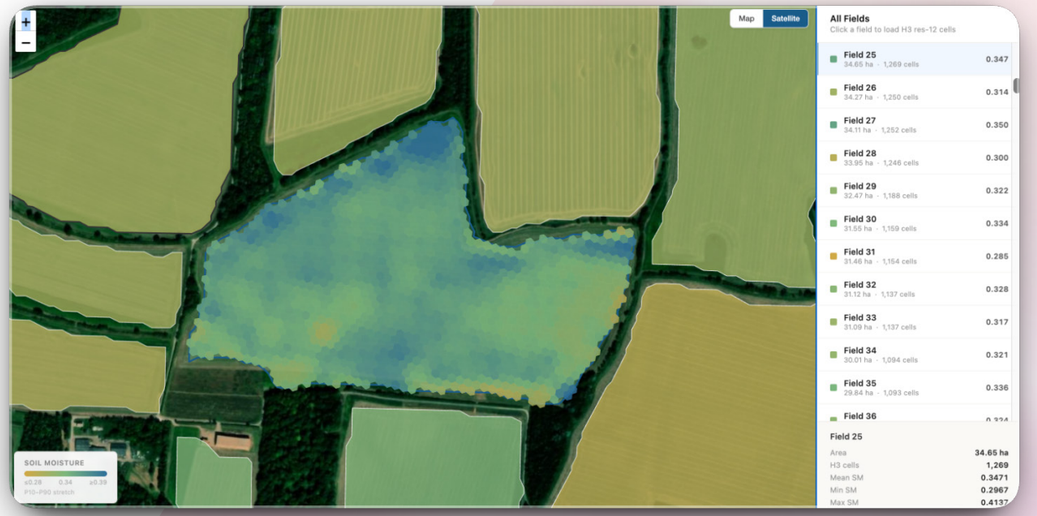
USE CASE: Field-scale soil moisture prediction

THEME:
Climate Resilience & Environmental Intelligence

Fig. 07 • Field-scale soil moisture view

ManyWorlds prototype showing predicted soil-moisture variation across an individual field, with H3 res-12 cells layered over satellite imagery and field-level summary metrics shown alongside.

Fig. 07



> "Farmers, insurers and land managers do not just need data. They need insights they can act on with confidence."

10 m resolution

5-30 cm soil depth

1,536 vCPU scale target

48 concurrent compute nodes

Case Study

GEOTERRAIMAGE

GeoChat Analytics

A climate resilience assistant for KwaZulu-Natal.

CHALLENGE

KwaZulu-Natal is regularly affected by drought, floods and rainfall variability. The data needed to understand these risks is often held in GIS tools, spreadsheets and technical reports, making it hard for planners, extension officers and disaster managers to use when making time-sensitive decisions.

APPROACH

GEOTERRAIMAGE built an agentic geospatial application that turns plain-language questions into validated spatial queries. The assistant works across climate, crop, socio-economic and hydrology datasets, then returns answers through maps, charts and streamed reasoning.

WHAT WAS BUILT

The team built a full-stack GenAI geospatial demo with a FastAPI backend, AWS Strands Agent, SvelteKit frontend, Leaflet map and a PostGIS database. The system exposes more than 29 tools for table discovery, schema inspection, safe spatial SQL and data retrieval.

RESULTS

The prototype can answer questions about district vulnerability, enumeration-area variation, crop distribution, climate projections, dry spells and food-security risk. It includes a six-layer evaluation framework and more than 44 test cases across the agent tools.

AWS IMPLEMENTATION

The application runs in an AWS Landing Zone provisioned by Cloudscaler, using ECS Fargate, ALB, Cognito, RDS PostgreSQL/PostGIS, S3, ECR, CodePipeline, CodeBuild, Secrets Manager and CloudWatch.

NEXT STEPS

Seek pilot partners, publish results, integrate the tool into existing products and explore commercial opportunities.

REGION: Africa • South Africa

STAGE: Mid • Demo-ready

USE CASE: Climate risk and livelihood-vulnerability analysis

THEME:
Climate Resilience & Environmental Intelligence

Fig. 08

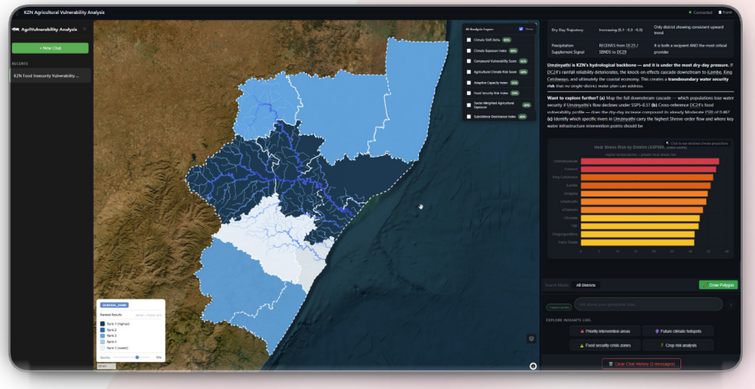


Fig. 08 • Food insecurity vulnerability clusters

GeoChat Analytics prototype showing agricultural vulnerability clusters across KwaZulu-Natal, highlighting food-insecure hotspot areas and district-level evidence for Umkhanyakude.

Fig. 09

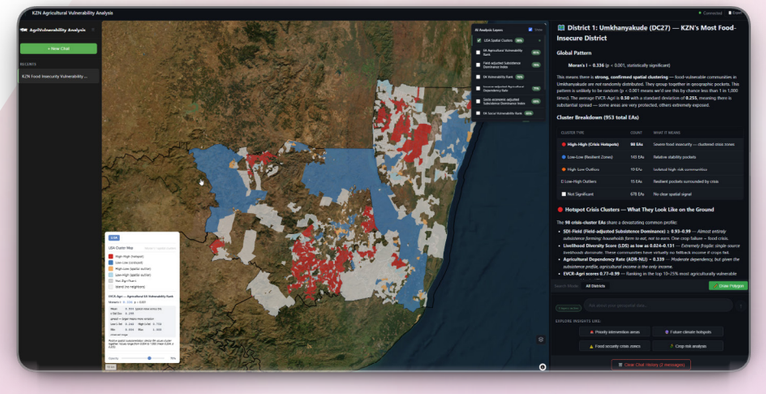


Fig. 09 • Climate and water-security risk view

GeoChat Analytics prototype showing climate and hydrological vulnerability layers for KwaZulu-Natal, combining district rankings with AI-generated analysis of dry-day trends, precipitation dependency and heat-stress risk.

> "In a province where climate shocks hit the poorest hardest, making geospatial climate intelligence accessible to non-experts is a step toward more equitable adaptation."

10 primary data sources

29 agent tools

44+ test cases

6-layer evaluation framework

Case Study

NatureAlpha

NATURE ALPHA

Ecosystem Service Forecasting

Forecasting ecosystem service risk with geospatial foundation models.

CHALLENGE

Nature-risk assessment is often backward-looking, coarse or incomplete. Companies and investors need to understand how water availability, habitat condition, vegetation stress and other ecosystem services may change around specific assets, not only across broad regions.

APPROACH

NatureAlpha used AlphaEarth foundation-model embeddings to analyse land-surface change over time. The team processed annual global data into basin-ready datasets, translated them into interpretable ecosystem-service layers and began forecasting how those services may shift over the next few years.

WHAT WAS BUILT

The team built a pipeline that turns satellite-derived embeddings into forward-looking maps of nature risk. The first build includes processing infrastructure, interpretable ecosystem layers and early forecasting models for ecosystem-service risk.

RESULTS

The pipeline is working and reproducible. Data processing is complete, an initial set of four ecosystem layers has been modelled from 2018 to 2030, and early forecasts are performing above chance, with further work planned to improve prediction quality.

AWS IMPLEMENTATION

Amazon S3 is used as the data lake and Amazon SageMaker for preprocessing, extraction, decoder training and forecasting, with Amazon Bedrock supporting LLM-based evaluation and output checks.

NEXT STEPS

Continue technical development, seek funding, publish results and integrate the approach into NatureAlpha's existing product and methodology stack.

REGION: Europe • United Kingdom

STAGE: Early • Prototype in progress

USE CASE: Asset-level ecosystem-service risk forecasting

THEME:
Climate Resilience & Environmental Intelligence

Fig. 10

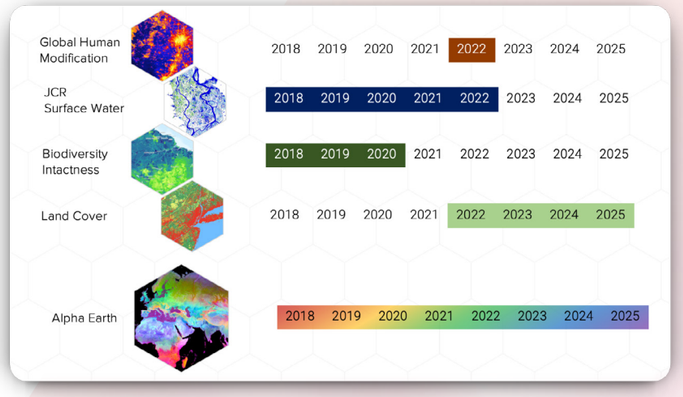


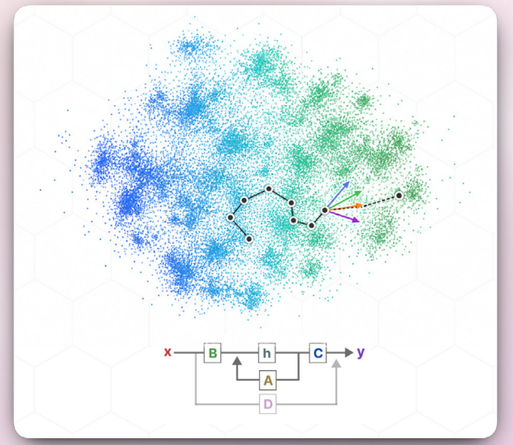
Fig. 10 • Ecosystem-service data timeline

NatureAlpha workflow showing how AlphaEarth embeddings align with global human modification, surface water, biodiversity intactness and land-cover datasets across different time periods.

Fig. 11 • Forecasting ecosystem-service change

NatureAlpha prototype visual showing how ecosystem-service signals move through the model to support early forecasting of future land-surface and biodiversity-risk shifts.

Fig. 11



> "Geospatial AI struggles with interpretability. By modelling geospatial layers across past and future, we can put novel geospatial AI methods to work to track critical ecosystem services."

26 dimensional embeddings

2017-2025 historical inputs

2018-2030 modelled layers

100 m forecast outputs

Case Study

Ocean OS

DARWIN & SeaLense

OCEAN OS

A marine geospatial foundation model and an agentic interface for offshore-wind consenting.

CHALLENGE

Offshore-wind consenting depends on fragmented marine data and specialist interpretation. First-pass marine-mammal scoping can take weeks, slowing site selection and increasing cost.

APPROACH

Ocean OS built two connected pillars: DARWIN, a marine geospatial foundation model, and SeaLense, an agentic offshore-wind assessment tool. The system combines RAG, Earth observation, vector geospatial data and cloud-native tools.

WHAT WAS BUILT

Ocean OS built two integrated pillars: DARWIN, a marine geospatial foundation model pretrained on 26 years of Copernicus ocean data, and SeaLense, a live web application that generates PEIR-style marine-mammal risk profiles for any sea area on demand.

RESULTS

DARWIN achieved 13.4% yearly mean relative error for chlorophyll-a prediction in 2022, outperforming baseline models. SeaLense is deployed live on AWS and matches a real Hornsea consenting assessment on data sources, reference populations and statutory regulatory flags.

AWS IMPLEMENTATION

DARWIN trained on Amazon EC2 p4d instances using Copernicus data staged in Amazon S3. SeaLense runs through Dockerised services using FastAPI, PostGIS and a STAC API, with the in-app agent served through Amazon Bedrock.

NEXT STEPS

Continue technical development, test with users, and pursue pilots with regulators, offshore-wind developers and environmental consultancies.

REGION:	Europe • United Kingdom
STAGE:	Mid • Working prototype
USE CASE:	Offshore-wind consenting

THEME:
Climate Resilience & Environmental Intelligence

Fig. 12

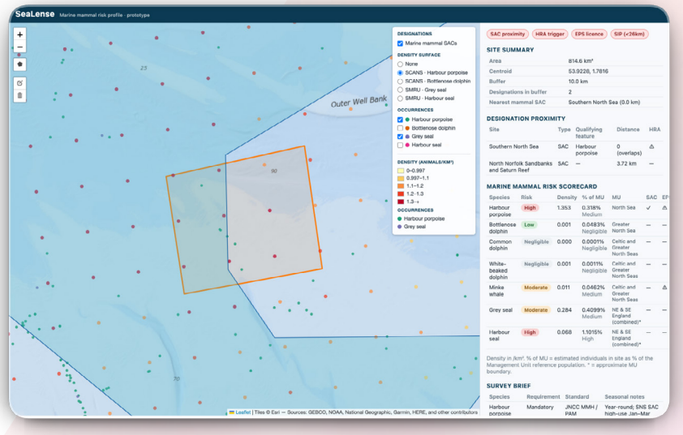


Fig. 12 • SeaLense prototype interface

North Sea site polygon with marine mammal occurrence data, density layers, protected area designations and risk indicators.

Fig. 13 • DARWIN global embeddings

January 2020 PCA visualisation of learned ocean-state representations from Copernicus Marine pretraining.

Fig. 14 • DARWIN European embeddings

Regional January 2020 PCA visualisation showing spatial patterns in DARWIN's learned marine data embeddings.

Fig. 13

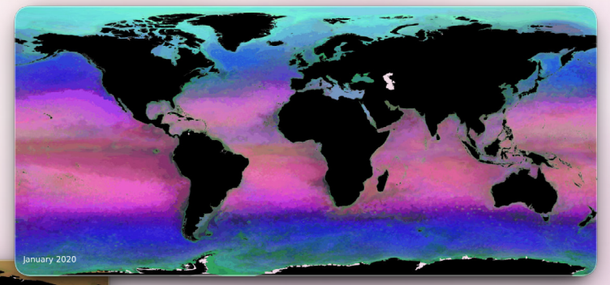
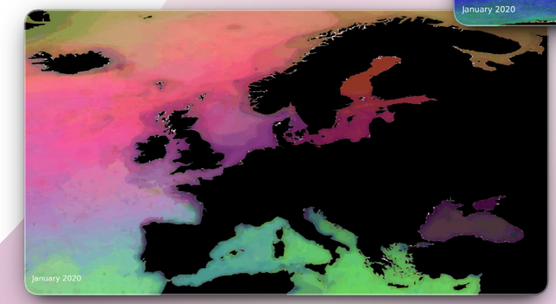


Fig. 14



> "The ocean is the least-understood part of our planet, yet we are building in it faster than ever."

26
years of Copernicus ocean data

13.4%
chlorophyll-a prediction error

10%
labelled data outperforming baseline

Weeks → seconds
marine-mammal scoping time

Case Study

Metaplanet

AURA

An agent-driven Earth observation platform for site-level risk assessment.

CHALLENGE

Insurance and risk-management teams need reliable site-level evidence on hazards, exposure and environmental change. Today, that work often depends on manual site visits, static hazard maps, fragmented GIS tools and specialist Earth observation analysis.

APPROACH

Metaplanet built AURA as an agent-driven platform that connects satellite data, TerraZard hazard layers, foundation-model outputs and natural-language interaction. The aim is not to replace risk experts, but to give them faster, traceable evidence for underwriting, claims and exposure analysis.

WHAT WAS BUILT

The team built an AWS-deployed demonstrator with a map interface, LLM geospatial agents, satellite AI models and TerraZard hazard layers. Users can ask about a site and receive maps, exposed-asset context, hazard history and an insurance-oriented explanation.

RESULTS

AURA has been demonstrated to insurance and risk-management stakeholders, including at the AMRAE conference. The team has developed a TerraZard flood demonstration, solved early cloud-deployment issues and is discussing proofs of concept with insurance companies.

AWS IMPLEMENTATION

AURA uses Bedrock-enabled LLM agents with AWS storage, compute, database, container and monitoring services. AWS credits supported deployment, experimentation and Bedrock testing.

NEXT STEPS

Continue technical development, improve the demo, seek pilot partners, seek funding, test with users, publish results, integrate into an existing product and explore commercial opportunities.

REGION: Europe • France

STAGE: Mid • Working prototype

USE CASE: Insurance underwriting, exposure and claims intelligence

THEME: Infrastructure & Smart Cities

Fig. 15 • Flood-depth hazard mapping

TerraZard interface showing flood-depth detection across Île-de-France, combining satellite imagery, water-depth estimates, permanent-water layers and built-area exposure context.

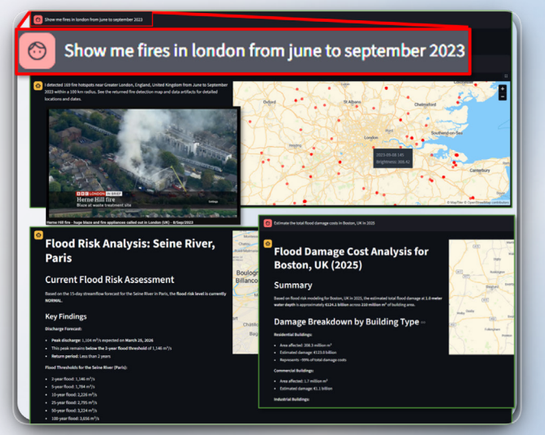
Fig. 15



Fig. 16 • Natural-language hazard queries

AURA demonstrator showing how users can ask site-specific risk questions and receive hazard maps, exposed-asset context and insurance-oriented analysis across fire and flood scenarios.

Fig. 16



> "AURA moves toward ready-to-wear intelligence: reusable AI-native representations, hazard layers and geospatial agents that embed expert knowledge into a service non-specialists can use."

AWS deployed demonstrator

TerraZard flood demo

10+ insurance demonstrators

PoCs in discussion

Case Study

VIDA

VIDA Agent

An agentic interface for infrastructure analytics.

CHALLENGE

Infrastructure and energy-access decisions often rely on large settlement and infrastructure datasets. VIDA already provides analysis-ready data across more than 60 countries, but complex ad hoc analysis can still require expert support, downloads or separate analytics workflows.

APPROACH

VIDA built an agentic layer that lets users ask questions of the map and the underlying data. The agent can query authorised workspaces, access data tools, interact with the map and return explanations through conversation.

WHAT WAS BUILT

The team built a Claude-on-Bedrock tool-use loop, a registry of tools for the agent and external MCP clients, and direct map interaction through panning, filtering and streamed frontend events. Every tool is scoped to the user's authorised workspace.

RESULTS

The VIDA agent is live in the platform, used internally by the data science team and actively demoed to customers. The team also ran an in-house hackathon to identify use cases and missing tools.

AWS IMPLEMENTATION

The agent uses Amazon Bedrock as the model layer, calling Claude through Bedrock cross-region inference profiles. The system uses prompt caching to reduce latency and cost, with Auth0 and Oso controlling authentication and workspace-level authorisation.

NEXT STEPS

Continue technical development and test with users. VIDA expects the agent to become a core part of its product offering over the next 12 months.

REGION: Europe • Germany

STAGE: Advanced • Being tested with users or partners

USE CASE: Infrastructure analytics and energy-access planning

THEME: Infrastructure & Smart Cities

Fig. 17

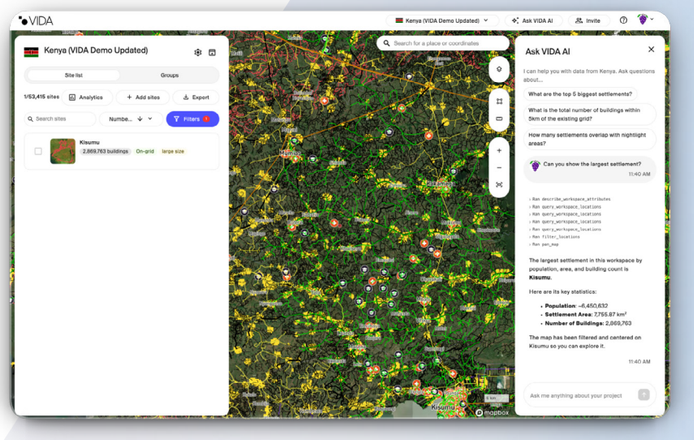


Fig. 17 • Nightlight coverage assessment

VIDA agent summarising nightlight coverage across Kenya, showing mapped settlements and AI-generated electrification insights.

Fig. 18 • High-opportunity site ranking

VIDA agent answering settlement queries in Kenya with mapped results on population, buildings and grid access.

Fig. 19 • Natural-language infrastructure query

VIDA agent ranking priority infrastructure sites by county to support grid-access planning.

Fig. 18

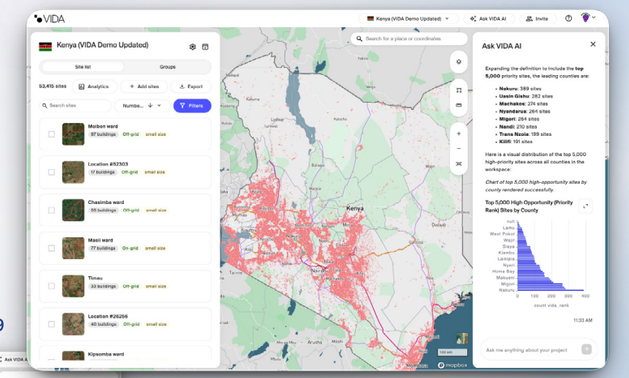
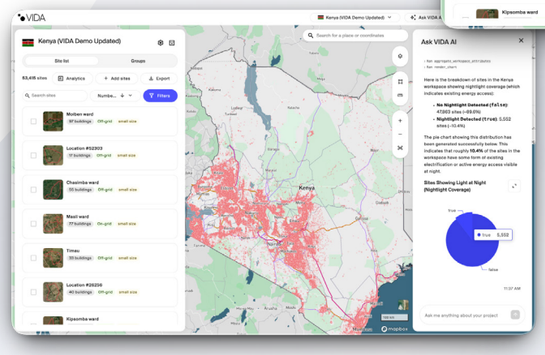


Fig. 19



> "From a map you read to an analyst you ask."

60+ countries covered

Live agentic interface

12-month product roadmap

53,415 Kenya demo sites

Case Study

VISTA

VISTA Agentic AI

An agentic geospatial AI platform for evidence-based spatial decision-making.

CHALLENGE

Organisations hold large amounts of spatial information across satellite imagery, maps, reports, regulations, administrative records and sensor data. These sources are often fragmented, and interpreting them requires specialist expertise and time.

APPROACH

VISTA combines Earth observation, vector data, documents and foundation models through an agentic reasoning layer. The system is designed to help users ask spatial questions, evaluate interventions and reason across multiple sources without moving raw datasets into the model context.

WHAT WAS BUILT

Terramira built an agentic geospatial AI system integrated into its existing platform. The agent works through a dedicated geospatial API layer with tools for spatial buffers, joins, overlays and distance calculations, plus asynchronous task execution for longer-running operations.

RESULTS

The prototype operates on national-scale datasets and has led to a follow-up R&D project funded by ESA Phi-Lab Netherlands with a Dutch university and municipality. That work will focus on uncertainty-aware agentic systems and native support for geospatial vector data in Earth observation foundation models.

AWS IMPLEMENTATION

VISTA is deployed as a cloud-native, container-based application on AWS using Amazon ECS, ELB, WAF, RDS, S3, ElastiCache, CloudWatch and Secrets Manager.

NEXT STEPS

Continue technical development, seek pilot partners and test with users, including Dutch municipalities and semi-public asset-management organisations.

REGION: Europe • The Netherlands

STAGE: Advanced • Pilot-ready

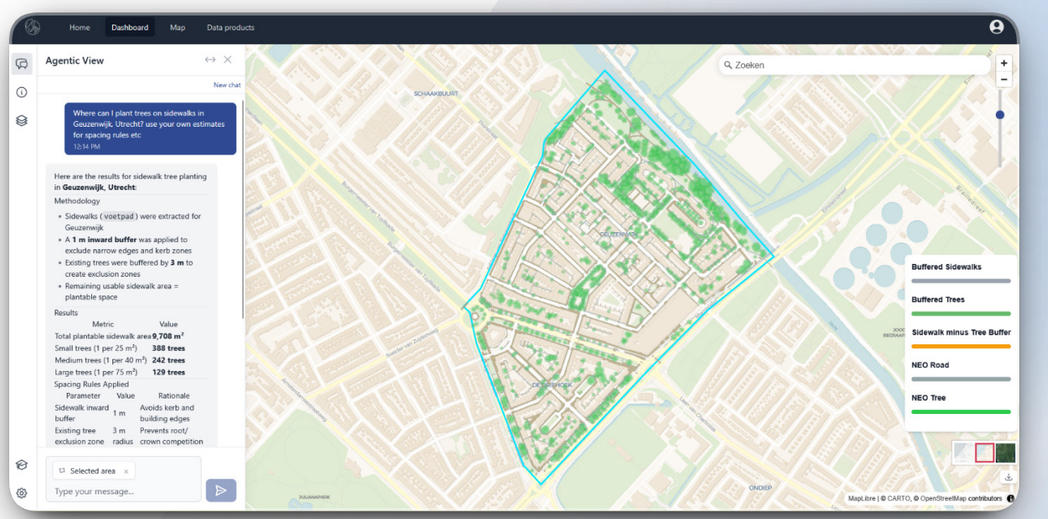
USE CASE: Planning, adaptation and environmental decision support

THEME: Infrastructure & Smart Cities

Fig. 20 • Sidewalk tree-planting suitability

VISTA interface showing an agentic query for tree-planting opportunities in Utrecht, combining sidewalk buffers, existing tree exclusion zones and spatial rules to estimate plantable area.

Fig. 20



> "Geospatial data provides the foundation for AI systems that are grounded in the real world."

National-scale
datasets

30 cm
satellite imagery planned

8 cm
aerial imagery planned

2
semi-public prospects

Case Study

Urban Climate MENA

Temperature Spatiotemporal Downscaling

Neighbourhood-scale urban heat intelligence pipeline.

CHALLENGE

Extreme urban heat is a growing climate and public-health risk across MENA. Current satellite products force a trade-off: Landsat offers 30 m detail but infrequent coverage, while MODIS provides frequent imagery at 1 km resolution. This makes neighbourhood-scale heat analysis difficult.

APPROACH

Urban Climate MENA used machine learning and multi-source satellite data to downscale MODIS land surface temperature to 30 m resolution. The model combines Landsat, MODIS, land cover and aerosol correction to deliver time-resolved urban heat intelligence.

WHAT WAS BUILT

The team built an end-to-end production-deployed pipeline on AWS and a live interactive dashboard. The system outputs error-corrected 30 m temperature for any time of day, including dates without a Landsat overpass.

RESULTS

The pipeline generates a 35-million-pixel scene in seconds and reaches around 1.5°C accuracy against held-out 2025 Landsat in its strongest months. The demo resolves neighbourhood-scale heat contrasts across Kuwait City and can scale to other cities.

AWS IMPLEMENTATION

Amazon S3 is used for large satellite training datasets, SageMaker training jobs and GPU instances for model development, DynamoDB for querying outputs by location and time, and Lambda for model and dashboard endpoints.

NEXT STEPS

Seek funding, test with users, publish results and explore commercial opportunities with health, meteorological and urban-planning agencies.

REGION: Middle East • Kuwait

STAGE: Advanced • Being tested with users or partners

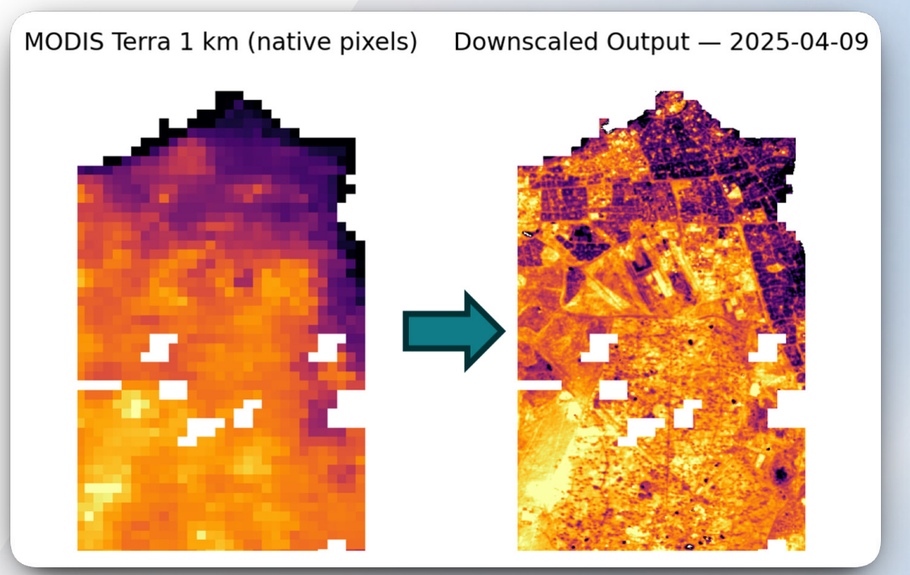
USE CASE: Urban heat monitoring and public-health planning

THEME: Infrastructure & Smart Cities

Fig. 21 • From 1 km MODIS to 30 m urban heat detail

Urban Climate MENA output comparing native 1 km MODIS land surface temperature with the downscaled 30 m result, revealing finer neighbourhood-scale heat patterns across Kuwait City.

Fig. 21



> "Our project matters because it brings city-scale heat intelligence to one of the world's most heat-stressed regions."

1 km → 30 m downscaling

35 million pixels in seconds

~1.5 °C strongest-month accuracy

25 years of satellite history

Case Study

ATOM AI

ATOM AI

ATOM AI Platform

A contextual AI data platform for geospatial and multimodal intelligence.

CHALLENGE

AI datasets often lack the geographic, cultural and local context needed for reliable decisions in underrepresented regions, while geospatial data can remain isolated as map features rather than being usable across sectors.

APPROACH

ATOM AI built a GenAI-powered platform that enriches raw multimodal data with geospatial and domain context, combining automated extraction, AI-assisted generation and local expert validation.

WHAT WAS BUILT

The team built an ATOM AI Platform prototype connected to its five-layer Data Engine: Ingestion, Curation, Contextualisation, Anonymisation and Distribution. It includes geospatial context mapping, configurable granularity, validation workflows, role-based interfaces and a visual orchestration canvas for AI data pipelines.

RESULTS

The prototype now connects backend APIs with a production-ready frontend and supports real-world AI data workflows. ATOM AI also completed three client data-curation projects and continued expanding its human-in-the-loop contributor network across African research and domain-expert communities.

AWS IMPLEMENTATION

AWS deployment uses Bedrock for GenAI context, Location Service for geospatial intelligence, Neptune for context graphs, Lambda and Step Functions for orchestration, and S3 for scalable storage, with EC2, RDS and related pipelines also in development.

NEXT STEPS

Continue technical development, seek pilot partners, pursue funding, test with users and expand enterprise pilots for contextual AI data preparation across Africa and emerging markets.

REGION: Africa • Kenya, DRC, South Africa, Tanzania, Rwanda

STAGE: Advanced • Being tested with users or partners

USE CASE: Contextual AI datasets and geospatial data enrichment



THEME:
Data Foundations, Trust & Scalable Platforms

Fig. 22

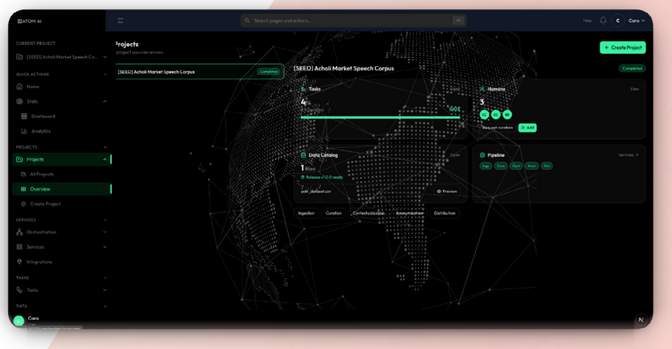


Fig. 22 • Project overview dashboard

ATOM AI Platform dashboard showing an active data-curation project, with task progress, human contributors, data catalogue status and the five-layer pipeline from ingestion to distribution.

Fig. 23 • Human-in-the-loop orchestration

Visual workflow showing how human validation connects with ingestion, curation, contextualisation, anonymisation and distribution.

Fig. 24 • Geospatial contextualisation view

Geospatial Contexts module showing mapped data points, cultural context and dataset rows.

Fig. 23

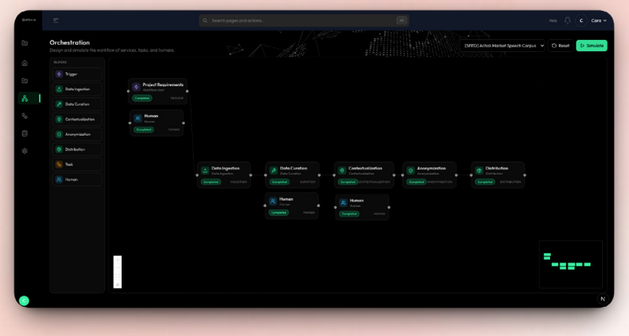
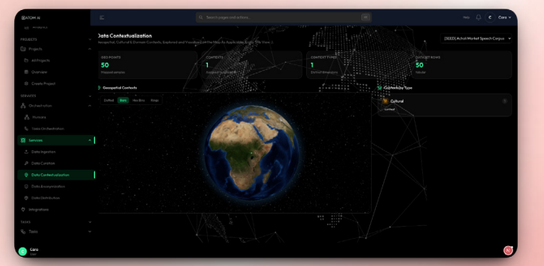


Fig. 24



> "Our project matters because future AI systems need more than data; they need context."

5-layer data engine

5 African countries represented

3 client data projects

Human-in-the-loop validation

Case Study

BigData Ghana

ForestTrace AI

An AI-assisted forest monitoring platform for faster environmental intelligence.

CHALLENGE

Forest degradation, illegal logging, illegal mining, wildfire risk, encroachment and land-use change are often detected too late for authorities and conservation teams to respond effectively. Existing monitoring can be slow, technical and difficult for non-specialists to interpret.

APPROACH

BigData Ghana built ForestTrace AI around satellite imagery, forest cover data, vegetation indicators, protected-area boundaries, environmental-risk layers and an AI interpretation layer. The focus is to make verified forest-monitoring outputs easier to question, summarise and use.

WHAT WAS BUILT

The team developed the geospatial data foundation, dashboard workflow, cloud architecture and forest-monitoring pipeline. The AI assistant will let users query verified forest loss, degradation, land-use change and risk outputs.

RESULTS

ForestTrace AI moved from a broad forest-monitoring concept to a structured prototype with a clear demo pathway. The project now has a defined data pipeline, dashboard direction and AI-ready workflow for forest intelligence.

AWS IMPLEMENTATION

The planned AWS architecture uses S3 for storage, compute for processing, RDS/PostgreSQL with PostGIS for spatial data, Lambda for automation, CloudWatch for monitoring, IAM for access control and Bedrock for GenAI.

NEXT STEPS

Continue development, prepare the demo, seek pilot partners, seek funding, test with users and explore commercial opportunities with forestry agencies, conservation NGOs, REDD+ actors, carbon projects and land-use organisations.

REGION: Africa • Ghana

STAGE: Early • Prototype in progress

USE CASE: Forest monitoring and deforestation-risk intelligence

THEME:
Data Foundations, Trust & Scalable Platforms

Fig. 25

Fig. 25 • Live forest monitoring dashboard

ForestTrace AI interface showing live monitoring across Ghana, with regional forest-loss rankings, fire hotspots, alerts and protected-area layers.

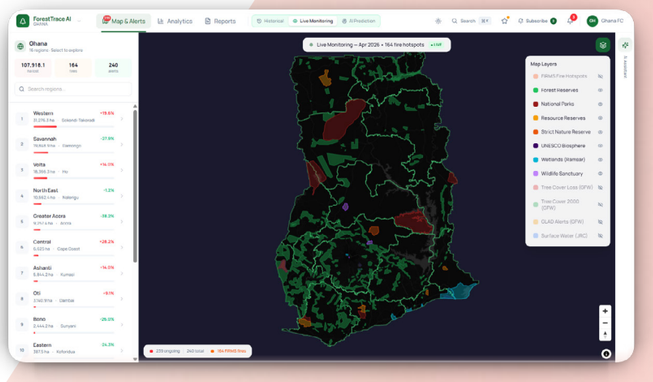


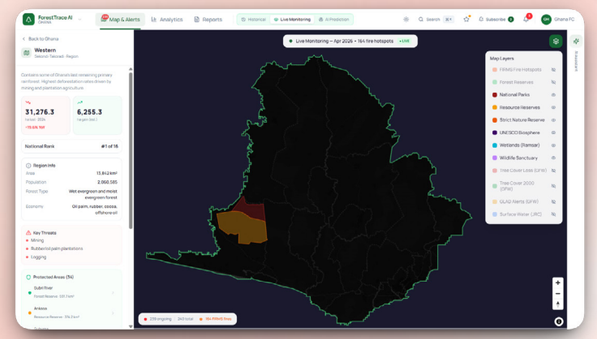
Fig. 26 • Regional forest-risk profile

Regional ForestTrace AI view showing forest-loss metrics, protected areas and key threats in Ghana's Western Region.

Fig. 26

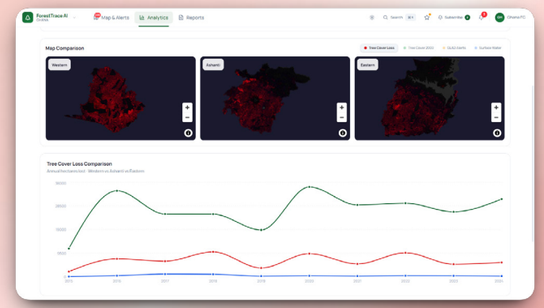
Fig. 27 • Natural-language infrastructure query

ForestTrace AI analytics view comparing regional tree-cover loss through maps and annual trend lines.



ForestTrace AI matters because faster forest intelligence can protect ecosystems before damage becomes irreversible.”

Fig. 27



16
Ghanaian regions monitored

24
years of forest-cover data

240
alerts tracked

African-built
forest intelligence

Case Study

RASID

GoPilot

An agentic Earth observation platform for methane monitoring and beyond.

CHALLENGE

Methane emissions from oil and gas infrastructure, landfills and agricultural operations can go undetected for long periods. Satellite imagery can help, but turning raw Earth observation data into reliable insight still requires specialist workflows.

APPROACH

RASID developed GoPilot, a geospatial agentic AI assistant for natural-language Earth observation analysis. It combines a purpose-built methane detection model with workflows that retrieve data, run tools and answer operational questions.

WHAT WAS BUILT

The team built GoPilot and MethaneMapper, a UNet-based methane plume detection model trained on physics-based synthetic plume data from Sentinel-2 imagery. GoPilot can also run NDVI and crop phenology analysis, change detection, land-cover classification and building-footprint retrieval.

RESULTS

MethaneMapper detected over 75% of real methane plume events in validation, despite being trained only on synthetic imagery. GoPilot demonstrated natural-language geospatial reasoning and was showcased at AWS Geospatial Day London on 2 June 2026, generating follow-up interest.

AWS IMPLEMENTATION

RASID used SageMaker Studio for data generation and model training, with datasets and model artefacts stored in private S3 buckets. GoPilot was built with Amazon Bedrock, AgentCore Runtime, Strands Agents and Claude for reasoning.

NEXT STEPS

Continue technical development, seek funding, test with users and explore commercial opportunities across environmental monitoring, agriculture, infrastructure, climate intelligence and wider Earth observation workflows.

REGION: Middle East • Lebanon

STAGE: Advanced • Being tested with users or partners

USE CASE: Methane monitoring and conversational EO analysis

THEME:
Data Foundations, Trust & Scalable Platforms

Fig. 28

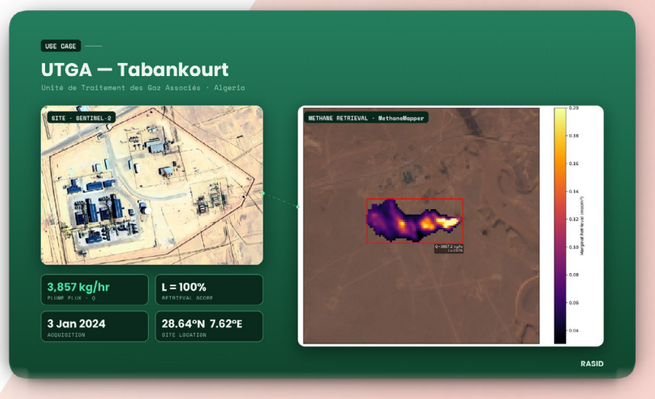


Fig. 28 • MethaneMapper detects methane plumes from Sentinel-2 imagery

Example methane plume detected at the UTGA gas treatment facility (Algeria). The AI model identifies plume extent and supports emission quantification from satellite imagery.

Fig. 29

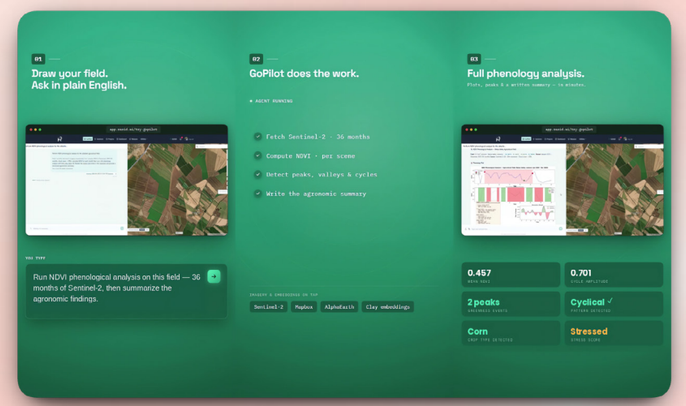


Fig. 29 • GoPilot workflow for natural-language crop phenology analysis

A user submits a natural-language request, GoPilot automatically retrieves Sentinel-2 imagery, computes NDVI, detects phenological patterns, and generates visualizations, metrics, and an agronomic summary.

> "Interacting with satellite intelligence should be as easy as having a conversation."

100k+ synthetic methane plumes

75%+ real plume detection

3,095 Sentinel-2 tiles

Natural-language EO workflows

Case Study

Sproxil

Project INSIGHT

Integrated national settlement intelligence for geographic health targeting

CHALLENGE

Subnational tailoring depends on identifying how malaria risk and access to care vary by place. Settlement type is central, but Nigeria lacks a single boundary dataset that clearly distinguishes urban, rural and underserved peri-urban communities, making a coordinate alone insufficient for health planning.

APPROACH

Sproxil built Project INSIGHT on the Global Human Settlement Layer Settlement Model, using the Degree of Urbanisation method to classify Nigerian locations by settlement stratum. A generative AI layer explains what each result may mean for malaria diagnostics and treatment planning.

WHAT WAS BUILT

The team built a cloud-native web app where users can click a map point or enter coordinates to return a settlement stratum: Urban Centre, Urban Cluster, Rural, Water or No Data. The map and classification engine are working and publicly accessible, with an interpretation layer in place and natural-language querying in preparation.

RESULTS

The prototype returns a stratum within seconds and can stratify around 80,000 georeferenced malaria survey responses by settlement type. An MoU with Nigeria's National Malaria Elimination Programme provides a route to support existing stratification and targeting workflows.

AWS IMPLEMENTATION

The prototype runs on a cloud-native architecture, with AWS credits supporting hosting, testing and development. The generative interpretation layer is provided by a large language model accessed as a service.

NEXT STEPS

Continue development, secure pilot partners and funding, test with users, and integrate Project INSIGHT into Sproxil's wider platform and malaria programme workflows.

REGION: Africa • Nigeria

STAGE: Mid • Working prototype

USE CASE: Settlement-based targeting for malaria interventions

THEME:
Data Foundations, Trust & Scalable Platforms

Fig. 30

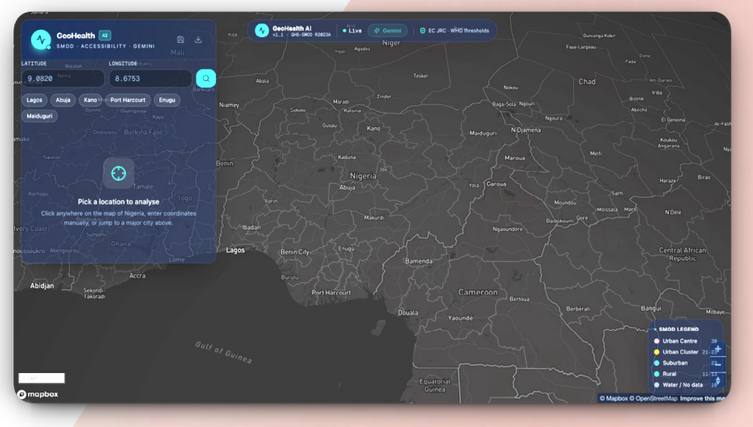


Fig. 30 • Project INSIGHT location analysis interface

Project INSIGHT AI map interface showing Nigeria-wide location selection for SMOD classification and health-accessibility analysis.

Fig. 31 • SMOD and health-accessibility results

Project INSIGHT AI output classifying a selected location as high-density urban and underserved for health access.

Fig. 32 • Nearest facility assessment

Project INSIGHT AI result showing the nearest health facility, travel distance and estimated access time for a selected urban location.

Fig. 31

Fig. 32

> "A coordinate alone cannot tell a health programme whether people live in a city or a remote village."

80,000 georeferenced malaria survey responses

4 settlement strata

Seconds to return a result

NMEP partnership route

Case Study

BolgiaTen

EO-VTI

A marine intelligence copilot for dark fleet risk.

CHALLENGE

AIS vessel tracking can be switched off, spoofed or jammed. This creates blind spots that can hide illegal fishing, sanctions evasion, illicit cargo transfers and underreported emissions. Regulators, insurers and maritime analysts need independent evidence of vessel behaviour.

APPROACH

BolgiaTen built EO-VTI to fuse satellite Earth observation imagery with AIS vessel tracking data. The system detects dark vessels, position spoofing and GPS jamming, and makes satellite-derived evidence easier to interpret through a natural-language layer.

WHAT WAS BUILT

The team built a proof-of-concept pipeline that ingests Sentinel-1 SAR, Sentinel-2 imagery and AIS data, fuses them spatially and temporally, and flags maritime anomalies. The platform uses Python, FastAPI, SQS workers, PostGIS and Bedrock-enabled GenAI features.

RESULTS

BolgiaTen demonstrated vessel detection on SAR and optical imagery, EO-AIS fusion over the Strait of Hormuz, and unique vessel tracking validated against hand-labelled test data. A customer workshop was delivered in April 2026, with a further demonstration to Lloyd's List Intelligence planned for July 2026.

AWS IMPLEMENTATION

EO-VTI is deployed in AWS eu-west-2 within a Cloudscaler-managed Landing Zone. It uses three Amazon ECS Fargate services, Amazon RDS PostgreSQL/PostGIS, Amazon SQS, Amazon S3, Amazon ECR, Amazon Bedrock and AWS Systems Manager Session Manager.

NEXT STEPS

Continue technical development, improve the demo, test with users and integrate into Lloyd's List Intelligence workflows, with a TRL-6 AWS prototype targeted for field testing by the end of September 2026.

REGION: Europe • United Kingdom

STAGE: Mid • Demo-ready

USE CASE: Maritime anomaly detection and vessel-risk intelligence

THEME:
Ecosystems, Markets & Decision Intelligence

Fig. 33

Fig. 33 • Dark-fleet intelligence workflow

BolgiaTen workflow showing how vessel detection, fingerprinting, correlation and risk assessment are combined through a unified data layer and language interface.

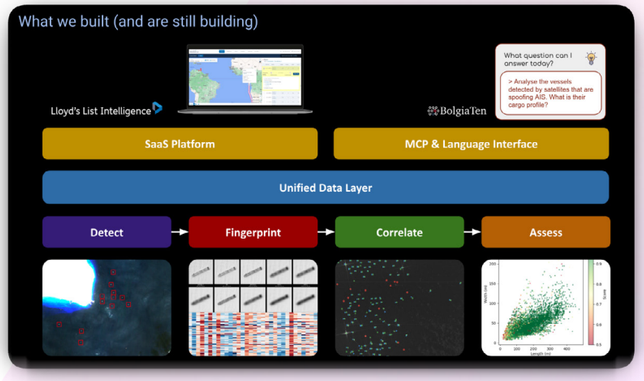


Fig. 34 • EO-enhanced Vessel Transparency Index

EO-VTI concept showing how multimodal vessel data feeds an intelligence engine to produce risk scores, alerts and explainable reports.

Fig. 34

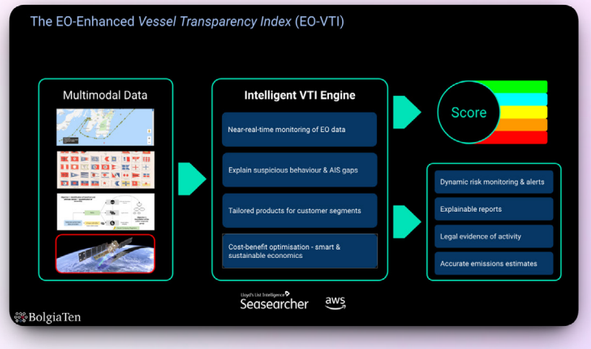
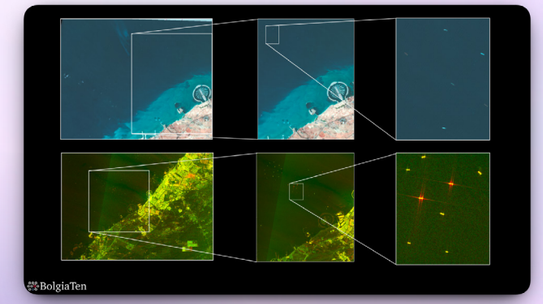


Fig. 35 • Satellite-based vessel detection

Satellite imagery examples showing vessel detections across optical and radar views, supporting hard evidence of maritime activity.

Fig. 35



> "AIS tells you where ships say they are, and satellite imagery tells you where they actually are."

3
core capabilities demonstrated

Radar
+ optical vessel detection

EO-AIS
fusion in Strait of Hormuz

TRL-6
target by September 2026

Case Study

Globeholder AI

Globeholder AI

UrbanRegion

Multimodal pipeline for region-level urban intelligence.

CHALLENGE

Many geospatial AI systems represent the world as pixels, tiles, points or boxes. Urban decisions need richer context: neighbourhoods, roads, POIs, street-level and satellite data, graph structure, text and trusted explanations.

APPROACH

Globeholder AI fuses multimodal physical and semantic signals to represent arbitrary regions. UrbanRegion builds city regions, learns region-level embeddings and enables agentic natural-language access.

WHAT WAS BUILT

The team built an urban intelligence pipeline that turns multimodal city data into validated region-level embeddings and agent-ready map outputs, combining neighbourhood boundaries, roads, POIs, Street View, satellite data, text, graph structure and location context.

RESULTS

The prototype demonstrates region-level similarity and interpretation across urban areas, helping analysts explore similar neighbourhoods, expansion opportunities, underserved areas and supporting evidence.

AWS IMPLEMENTATION

Globeholder AI uses Amazon S3 for the batch processing layer and Amazon Bedrock for the agent layer, supporting multimodal data processing and natural-language access to region-level urban intelligence outputs.

NEXT STEPS

Continue developing Globeholder AI as an urban-region intelligence layer, strengthening region generation, multimodal fusion, validation and natural-language access.

REGION: Europe • France

STAGE: Mid • Working prototype

USE CASE: Urban-region intelligence and neighbourhood comparison

THEME:
Ecosystems, Markets & Decision Intelligence

Fig. 36

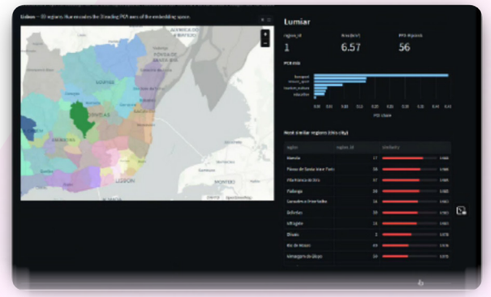


Fig. 36 • Region embedding explorer

Globeholder AI interface showing region-level embeddings for Lisbon, with neighbourhood similarity scores, POI breakdowns and map-based regional comparison.

Fig. 37

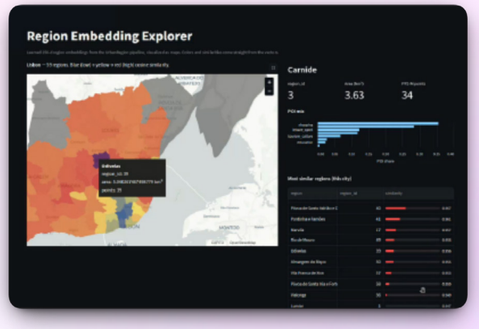


Fig. 37 • Cross-region similarity analysis

Region Embedding Explorer output for Carnide, showing how the system identifies similar urban regions using embeddings, POIs and spatial context.

Fig. 38 • Similar-region ranking

Similarity ranking from the Globeholder AI prototype, comparing neighbourhoods within the same city and across other cities from the same model run.

Fig. 38



> "Urban decisions need explainable regional intelligence."

59 Lisbon regions

7-step intelligence pipeline

8 multimodal signal types

Region-level embeddings

Key Takeaways

Across the cohort, one message came through clearly: generative AI can make geospatial workflows easier to access and use, but only when paired with strong data foundations, spatial expertise, responsible design and flexible cloud infrastructure. AWS support gave teams the freedom to work across different models and increasingly multimodal workflows, combining text, imagery, geospatial data and other sources as they moved from generative AI experimentation towards more agentic, operational systems.

Seven stood out



01

AI works best as an interface, not a shortcut



Many teams used generative AI as an interface, reasoning layer or workflow assistant, rather than a replacement for geospatial expertise. The strongest prototypes used AI to help users ask better questions, retrieve data, run tools, summarise outputs and understand what a result means.

02

Analysis readiness remains the hidden work



Before models could generate useful outputs, teams had to prepare, clean, structure and validate their underlying geospatial data. This included raster and vector data, satellite imagery, field data, public datasets, hazard layers, environmental indicators and synthetic training data. The quality of each prototype's output depended heavily on the quality of the data pipeline behind it.

03

The required skill set is broadening



The challenge showed that geospatial AI calls for a wider mix of skills than traditional GIS or AI alone. Teams worked across cloud-native data, agentic workflows, model evaluation, geospatial methods, cost management, user experience and communication. The upskilling need has not disappeared – it has moved up the stack.

04

Non-technical access still needs translation



Natural-language interfaces can lower the barrier to using geospatial tools, but they do not remove the need for explanation. Users still need to understand where an answer came from, what assumptions were made, what uncertainty remains, and what decisions the output can or cannot support.

05

Infrastructure works best with support around it



AWS credits, cloud infrastructure and technical support gave teams space to experiment, train models, run heavier processing jobs and deploy prototypes. The Build Phase added another layer of value by giving teams structured check-ins, technical Q&A, Slack support and peer visibility while they worked through practical implementation challenges.

06

Responsible AI is a practical discipline, not an abstract one



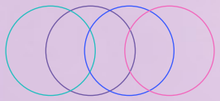
For these projects, responsibility was not only a policy question – it showed up in practical choices around data quality, compute cost, permissions, uncertainty, validation and human oversight. The strongest projects were clear-eyed about what their systems could support, and where specialist review still mattered.

07

A prototype is not the same as adoption



By the end of the Challenge, teams had reached different stages of maturity – from early prototype to demo-ready, pilot-ready or already being tested with users. The next challenge is adoption: proving who needs the tool, what decision it supports, and what evidence would give users enough confidence to act.



What's Next

The challenge ends with a strong set of prototypes, but the next stage is where the hardest questions begin.

The Challenge ends with a strong set of prototypes, but the next stage is where the hardest questions begin.

Teams are now moving in different directions, from pilots and user testing to product integration, funding conversations, research publication and commercial development. Some still need further technical refinement before they can be tested in real operational settings.

For many, the challenge does not end with a prototype. Their work is now progressing towards pilot deployments, operational services and commercial partnerships. This reflects a wider industry shift from experimenting with generative AI towards deploying agentic AI systems that can orchestrate workflows, reason across multiple data sources and support real operational decision-making.

The next phase will not look the same for every project. Some teams need pilot partners; others need stronger cloud architecture or AI and machine learning support, clearer routes to public-sector adoption, procurement or policy use, or help communicating their value to users who may not speak the language of geospatial AI.

Across the cohort, the shared next step is to move from demonstration to evidence of use. That means testing whether each tool answers the right question, fits into existing workflows, works within real-world data constraints and gives users enough confidence to act.

The challenge created the conditions for that next step: time, credits, technical support, visibility and a cohort to build alongside. Its value lies not only in the prototypes produced, but in the learning generated around what practical geospatial AI now requires.

This is not the end of the work – it is the point at which the strongest ideas begin to meet the users, institutions and decisions they were built to support.

Different routes forward

- Pilot partners
- User testing
- Research publication
- Cloud architecture
- Product integration
- Value communication
- Commercial development
- Public-sector adoption
- Funding

> "This challenge was never just about building prototypes in five months. It was about proving that generative AI and geospatial data, brought together with the right support, can move from idea to something real."

– Julia Wagemann, Co-CEO thriveGEO

Acknowledgements

The GenAI for Geospatial Challenge – EMEA Edition was made possible through the collaboration, support and participation of a wide community.

thriveGEO would like to thank **AWS** for co-hosting the challenge, and **Women+ in Geospatial**, and **LunateAI** for their support in helping shape and strengthen this edition. Thanks also to **Cloudscaler** for their landing zone support for our winning teams.

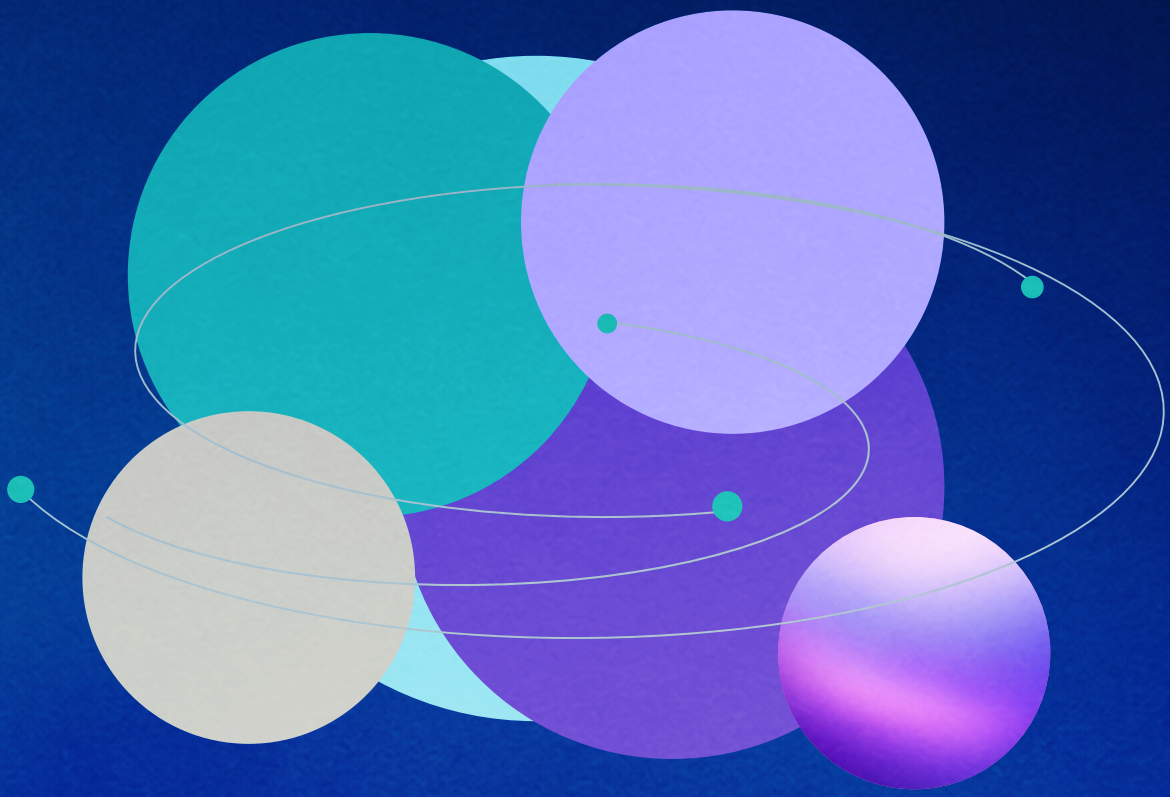
We are also grateful to the judging committee, mentors, speakers and technical contributors who gave their time, expertise and feedback throughout the challenge. Their input helped teams sharpen their ideas, navigate technical questions and think more clearly about the real-world relevance of their solutions.

We would also like to recognise the speakers, moderators and panellists who contributed to the GenAI for Geospatial Challenge Kick-off event in London: Rebecca Meads (AWS); Phil Cooper (AWS); Javier de la Torre (CARTO); Samuel Barrett (LGND AI); Nadine Alameh (LunateAI); James Banting (Sparkgeo); Josué Molina (LuxCarta); Emily Burke (AWS); Ryan Averay (Wherobots); Donna Lyndsey (Donna Lyndsey Consulting); Shikoh Gitau (Qhala); Ed Parsons (Digital Geographer); Kat Semperi (ICEYE); Ahmad Aldhalaan (GeoTech); Eren Unlu (Globeholder); Andrew Pawloski (Element84); and Lucas Bishop (Seraphim Space).

Most importantly, we would like to thank every team that applied. The quality, ambition and range of submissions showed the depth of creativity already emerging at the intersection of generative AI and geospatial innovation across Europe, the Middle East and Africa.

To the winning teams: thank you for bringing energy, openness and commitment to the Build Phase. Your prototypes, questions and lessons helped turn the challenge into more than a competition. They made it a shared space for experimentation, learning and practical progress.

This report reflects the work of that wider community, and we are grateful to everyone who contributed to it.



Don’t just observe the future of geospatial AI – build it.

Join us. Innovate with us. Lead with us.

GenAI for Geospatial Challenge

EMEA Edition 2025-2026

Hosted by thriveGEO and AWS, with the support from Women+ in Geospatial and LunateAI.