

AF-S1-100453

Resilia Pulse: Alignment-Aware Compliance AI for Safe and Transparent Financial Systems

The accelerating deployment of artificial intelligence in critical infrastructure presents both opportunities and risks. Gas infrastructure in Sub-Saharan Africa suffers from high downtime, leaks, and emergency flaring, costing operators billions annually and contributing to climate change. While predictive maintenance solutions exist, most are expensive, reactive, or lack alignment safeguards. Current systems often fail to detect harmful anomalies in time or overreact with false alarms that disrupt operations. Both outcomes reflect misaligned AI behavior in safety-critical domains.

Resilia Pulse is an AI-powered digital twin platform designed to ensure that AI systems monitoring energy assets remain safe, transparent, and under human control. The system integrates low-cost IoT sensors, cloud-based data flows, and predictive models to simulate asset health, forecast failures, and generate ESG compliance reports. Unlike conventional monitoring tools, Resilia Pulse explicitly embeds alignment principles into its architecture: human-in-the-loop oversight, explainability, and adversarial robustness.

Our approach consists of three pillars:

Alignment by Design – Predictive AI models will be constrained with safety and fairness parameters, ensuring they do not attempt harmful actions such as ignoring gas leaks, misreporting emissions, or prematurely shutting down critical systems.

Human-in-the-Loop Control – Operators will receive interpretable alerts and retain override authority. Audit trails and transparency modules will reinforce accountability and prevent over-automation.

Adversarial Testing at Scale – With dedicated compute resources from AWS, we will simulate attack scenarios, stress-test anomaly detection, and validate system reliability under edge cases that current literature rarely addresses.

This research is distinct from prior work, which focuses narrowly on either predictive maintenance or AI safety in abstract domains. By positioning energy infrastructure as a testbed for applied alignment, Resilia Pulse bridges a gap between theory and practice. Our project demonstrates how alignment safeguards can be embedded into industrial digital twins—an area often overlooked in alignment research.

The project aligns directly with The Alignment Project’s priority question: “How can we prevent AI systems from carrying out actions that pose risks to our collective security, even when they are attempting to carry out such actions?” By focusing on real-world risks—gas leaks, unreported emissions, financial loss, and safety hazards—Resilia Pulse reframes alignment as both a technical and societal necessity.

Expected outcomes include:

An alignment-aware compliance and monitoring prototype for gas infrastructure.

Open datasets of adversarial test cases for industrial AI.

Publications on alignment methodologies for digital twins.

Policy and ESG insights relevant to emerging markets.

By situating alignment research in the Global South and addressing energy security alongside AI safety, Resilia Pulse expands the scope of the alignment field, ensuring its benefits are globally inclusive and practically relevant.

CONTACT DETAILS

Name Johnpaul
Surname Okeke
Tel (Work) +2349032210788
(tel:+2349032210788)
Email (Work) okekejohnpaul12@gmail.com
(mailto:okekejohnpaul12@gmail.com)
Address

AF-S1-100453

Resilia Pulse: Alignment-Aware Compliance AI for Safe and Transparent Financial Systems

Section 1 - Eligibility Statement

To be eligible for the Alignment Project, I confirm that:

Checked

Section 2 - Due Diligence (Your Organisation)

Type of Organisation

Individual

Companies House Number (or equivalent)

Not yet registered – Resilia Pulse is in the process of incorporation under CAC, Nigeria

Amount of Funding Required

£200,000.00

Location where the funding will be spent

Nigeria – primary research, development, and testing (University of Lagos, Lagos).

Global – AWS cloud compute resources for experimentation.

Section 3 - Research Project Details

CONTACT DETAILS

Name	Johnpaul
Surname	Okeke
Tel (Work)	+2349032210788
Email (Work)	okekejohnpaul12@gmail.com
Address	

Contact Title/Role

Founder & Lead Researcher, Resilia Pulse

GMS ORGANISATION

Type	Organisation
Name	Independent (Resilia Pulse Project)
Phone (Work)	+2349032210788
Email (Work)	okekejohnpaul12@gmail.com
Address	University of Lagos, Lagos, Nigeria LAGOS Nigeria

Affiliated Institution Country Location

Nigeria

Relevant Regulatory Authority

N/A (Independent Researcher)

Legal Registration Number

N/A

Orcid ID / Google Scholar / LinkedIn Profile Link

<https://www.linkedin.com/in/okeke-johnpaul-659685173/>

Additional Collaborators

Temitope Jinadu – IoT Engineer, University of Lagos (Co-developer of Resilia Pulse hardware prototype).

Dr Olugbenga Olamigoke – Senior Lecturer, Petroleum and Gas Engineering, University of Lagos (Academic mentor, research guidance).

Sirius-X Energy (Zero Emissions Project) – Field Research Support (Data collection and validation collaboration).

Relevant Research Outputs

Conference Paper – SPE NAICE 2025

Programmable Nanoparticles for Smart Drilling Mud Systems in Over-Pressured Shale Formations

Accepted for presentation at the Society of Petroleum Engineers (SPE) Nigeria Annual International Conference & Exhibition (NAICE) 2025.

Focuses on nanotechnology applications in drilling systems, with potential cross-links to AI-driven predictive models.

Technical Project – EnerDeFi IDHub (2025)

AI-driven Compliance & Identity Verification System

Developed a multi-API KYC/AML compliance engine integrating Mono, Preamble, and VerifyMe, with features such as audit trails, sanctions/PEP screening, and regulatory reporting.

Relevant to alignment research in financial AI transparency and human control.

Educational AI Platform – CBT Grinder (2025)

AI-powered Adaptive Learning Platform for High-Stakes Exams

Co-founded CBT Grinder, a multi-tenant exam prep platform featuring AI tutors, offline-first access, and predictive performance analytics. Demonstrates applied AI in education, explainability, and fairness.

Research Assistant Role – Advisory Reports (2024–Present)

Published Energy & Sustainability Reports

Contributed to policy and market analysis in energy transitions, leveraging data science and AI. Published outputs on AI's role in sustainability and resilience.

Prototype – Resilia Pulse (2024–2025)

AI-Powered Digital Twin for Predictive Gas Asset Monitoring

Built a functional prototype using IoT sensors + AI digital twin models to monitor gas infrastructure health, predict failures, and reduce downtime/leaks.

Demonstrated live anomaly detection and ESG reporting in pilot simulations

Research Title

Resilia Pulse: Alignment-Aware Compliance AI for Safe and Transparent Financial Systems

Research Statement

The accelerating deployment of artificial intelligence in critical infrastructure presents both opportunities and risks. Gas infrastructure in Sub-Saharan Africa suffers from high downtime, leaks, and emergency flaring, costing operators billions annually and contributing to climate change. While predictive maintenance solutions exist, most are expensive, reactive, or lack alignment safeguards. Current systems often fail to detect harmful anomalies in time or overreact with false alarms that disrupt operations. Both outcomes reflect misaligned AI behavior in safety-critical domains.

Resilia Pulse is an AI-powered digital twin platform designed to ensure that AI systems monitoring energy assets remain safe, transparent, and under human control. The system integrates low-cost IoT sensors, cloud-based data flows, and predictive models to simulate asset health, forecast failures, and generate ESG compliance reports. Unlike conventional monitoring tools, Resilia Pulse explicitly embeds alignment principles into its architecture: human-in-the-loop oversight, explainability, and adversarial robustness.

Our approach consists of three pillars:

Alignment by Design – Predictive AI models will be constrained with safety and fairness parameters, ensuring they do not attempt harmful actions such as ignoring gas leaks, misreporting emissions, or prematurely shutting down critical systems.

Human-in-the-Loop Control – Operators will receive interpretable alerts and retain override authority. Audit trails and transparency modules will reinforce accountability and prevent over-automation.

Adversarial Testing at Scale – With dedicated compute resources from AWS, we will simulate attack scenarios, stress-test anomaly detection, and validate system reliability under edge cases that current literature rarely addresses.

This research is distinct from prior work, which focuses narrowly on either predictive maintenance or AI safety in abstract domains. By positioning energy infrastructure as a testbed for applied alignment, Resilia Pulse bridges a gap between theory and practice. Our project demonstrates how alignment safeguards can be embedded into industrial digital twins—an area often overlooked in alignment research.

The project aligns directly with The Alignment Project's priority question: "How can we prevent AI systems from carrying out actions that pose risks to our collective security, even when they are attempting to carry out such actions?" By focusing on real-world risks—gas leaks, unreported emissions, financial loss, and safety hazards—Resilia Pulse reframes alignment as both a technical and societal necessity.

Expected outcomes include:

An alignment-aware compliance and monitoring prototype for gas infrastructure.

Open datasets of adversarial test cases for industrial AI.

Publications on alignment methodologies for digital twins.

Policy and ESG insights relevant to emerging markets.

By situating alignment research in the Global South and addressing energy security alongside AI safety, Resilia Pulse expands the scope of the alignment field, ensuring its benefits are globally inclusive and practically relevant.

Research Priority Areas

Empirical investigations into AI monitoring and red teaming

Deliverables and Timeline

The Resilia Pulse project will follow a structured timeline to ensure technical development, alignment validation, and dissemination of results. The work plan is divided into four key phases with concrete deliverables at each stage.

Phase 1 – Research & Framework Development (Months 1–2)

Conduct literature review on AI alignment in critical infrastructure and predictive maintenance.

Define alignment-specific requirements for gas infrastructure monitoring, focusing on adversarial robustness and human oversight.

Curate synthetic datasets from IoT sensors (pressure, temperature, gas leak) for baseline model training.

Deliverables:

Alignment design framework for predictive industrial AI.

Dataset of initial gas infrastructure test cases.

Phase 2 – Prototype Build & Early Testing (Months 3–4)

Develop the Resilia Pulse prototype integrating IoT sensors, cloud sync, and AI digital twin models.

Implement initial explainability features and operator override modules.

Run early-stage testing with synthetic anomalies (e.g., pressure drops, leak simulations).

Deliverables:

3. Prototype version 1.0 with alignment-aware monitoring functions.

4. Technical report on early-stage testing results.

Phase 3 – Adversarial Simulation & Human-in-the-Loop Integration (Months 5–8)

Leverage AWS compute resources to simulate adversarial and edge-case scenarios (fraudulent signals, sensor failures, high-flare events).

Refine predictive models with constraints that prevent harmful misclassifications.

Deploy human-in-the-loop interface for real-time decision override and explainability logs.

Deliverables:

5. Robust AI compliance engine tested against adversarial cases.
6. Integrated operator dashboard with explainability and override capability.
7. Public dataset of adversarial test cases for alignment research.

Phase 4 – Field Validation, Dissemination & Policy Impact (Months 9–12)

Conduct pilot deployment with partner organisations or synthetic field environments.

Evaluate alignment performance under real-world operating conditions.

Prepare and submit findings to peer-reviewed venues and AI alignment workshops.

Draft policy-oriented whitepaper linking industrial AI safety to ESG and global alignment frameworks.

Deliverables:

8. Field-validated prototype with alignment guarantees.
9. Research paper on alignment in digital twin systems.
10. Policy whitepaper on alignment-aware compliance AI.
11. Open-source components and documentation to support replication.

Overall Impact by Month 12

A fully tested prototype of Resilia Pulse, demonstrating alignment-aware AI for gas infrastructure.

Contribution of datasets, methodologies, and publications to the global alignment community.

Policy recommendations to ensure safe, auditable AI deployment in energy-critical domains.

Budget Narrative

Budget Narrative

We are requesting £200,000 to support a 12-month research program that develops, tests, and validates Resilia Pulse as an alignment-aware digital twin for gas infrastructure. The proposed budget is structured to ensure efficient use of resources while delivering high-impact technical and policy outcomes.

1. Personnel – £70,000

Research Lead (Student Founder Stipend): Supports time commitment for full-time coordination, model development, and dissemination.

IoT/Hardware Engineer (Part-time): Hardware integration, sensor calibration, and system maintenance.

Research Assistant/Data Analyst (Part-time): Data preparation, anomaly simulation, and ESG analytics. This allocation ensures the core team can dedicate time to rigorous technical development and alignment testing.

2. Compute & Technical Resources – £50,000

AWS Credits: Large-scale simulations for adversarial testing and model refinement.

Cloud Hosting & Data Storage: Secure handling of sensor streams, dashboards, and test case repositories.

Software Licenses/Development Tools: For ML frameworks, ESG compliance modules, and monitoring dashboards.

This is essential for training predictive models at scale and validating alignment guarantees under stress conditions.

3. Research Operations – £30,000

Prototype Hardware: IoT sensors (pressure, temperature, gas leak), controllers, and communications modules.

Field Testing: Controlled pilot deployments and validation with synthetic datasets.

Data Acquisition: Sourcing additional industrial datasets to enrich adversarial test scenarios.

These resources will demonstrate the feasibility of alignment-aware predictive monitoring in realistic environments.

4. Collaborations & Knowledge Sharing – £20,000

Workshops with academic mentors, compliance experts, and alignment researchers.

Collaboration with Nigerian energy stakeholders and international alignment partners.

This ensures cross-disciplinary input and broader impact beyond technical development.

5. Dissemination & Policy Impact – £15,000

Open-access publications in alignment and AI safety venues.

Whitepaper development on policy implications for safe AI in energy infrastructure.

Presentation at international AI safety and energy conferences.

6. Contingency – £15,000

Allocated to cover unexpected costs, such as hardware replacements or expanded testing requirements.

Total Request: £200,000

This budget strikes a balance between technical ambition and practical deliverability. It prioritises personnel, compute, and hardware testing, while also dedicating resources to dissemination and collaboration. By focusing on alignment in a high-stakes real-world domain, the investment will yield both applied technical outputs and contributions to the global alignment research agenda.

Additional information

As a student researcher and emerging founder, I bring a unique perspective to the alignment research community—bridging academic training in petroleum and gas engineering with practical experience in software engineering, AI development, and compliance systems. I believe this intersection of domains is essential for addressing the complex, real-world challenges of alignment.

My work with Resilia Pulse builds on a track record of projects at the nexus of AI, energy, and compliance. I co-authored a technical paper presented at the Society of Petroleum Engineers (SPE) Nigeria Annual International Conference & Exhibition (NAICE) 2025, led the development of EnerDeFi IDHub, an AI-driven identity and compliance platform, and co-founded CBT Grinder, an AI-powered education platform that supports thousands of learners. These projects demonstrate my ability to move ideas from research to prototype and into impact-oriented solutions.

What makes Resilia Pulse particularly relevant is its grounding in the Global South context. Many alignment projects focus on general-purpose AI models developed in Western labs, but fewer address the risks of AI systems in emerging markets where infrastructure is fragile, regulatory systems are evolving, and failures can have severe human and environmental consequences. By situating this work in Nigeria and Sub-Saharan Africa, Resilia Pulse ensures that alignment research reflects diverse contexts, contributing to a more inclusive global conversation.

The proposed research will also create open resources for the wider alignment community, including a dataset of adversarial test cases for industrial AI and a policy whitepaper linking AI safety to ESG reporting and compliance frameworks. These outputs are designed to benefit both academic researchers and industry practitioners.

Finally, my ambition is not only to develop technical safeguards but also to inspire more students and researchers in Africa to engage with AI safety and alignment. By demonstrating that alignment research can be both globally significant and locally impactful, I aim to expand the scope of participation in this critical field.

Resilia Pulse is therefore more than a technical project—it is an opportunity to broaden the alignment research ecosystem, bringing in new perspectives, addressing urgent safety challenges, and building systems that are both globally beneficial and locally resilient.

Section 4 - Applicant Declaration

You must read and agree to the stated declaration before submitting your application:

Checked