International Journal of Creative and Open Research in Engineering and Management
ISSN: 3108-1754 (Online)
Volume 02 Issue 11 November-2026 | Impact Factor: 3.5



# Aeris Ops: AI-Powered Smart Weather Intelligence, Disaster Prediction and Climate Analytics Platform

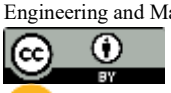
**Puja Gadhave[1], Shrutika Burande[2], Shilpa Khade[3], Prof. Komal Gaikwad[4], Dr. Bharti Sayankar[5], Dr. Sudhir N. Shelke[6]**

[1] Department of CSE, Guru Nanak Institute of Technology, Nagpur, India
[2] Department of CSE, Guru Nanak Institute of Technology, Nagpur, India
[3] Department of CSE, Guru Nanak Institute of Technology, Nagpur, India

[4] Department of CSE, GNIT (Project Guide), Nagpur, India
[5] Department of CSE, GNIT (Head of Department), Nagpur, India
[6] Principal & Director of (GNIT), Nagpur, India

**How to Cite this Article**:
Gadhave, P., Burande, S., Khade, S., Gaikwad, K., Sayankar, B. & Shelke, S. N. (2026). Aeris Ops: AI-Powered Smart Weather Intelligence, Disaster Prediction and Climate Analytics Platform. International Journal of Creative and Open Research in Engineering and Management, 02(11), 1-9. https://doi.org/10.55041/ijcope.v2i11.085

**License:**
This article is published under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.
© The Author(s). Published by International Journal of Creative and Open Research in Engineering and Management.





https://doi.org/10.55041/ijcope.v2i11.085

**Abstract**— Weather intelligence becomes operationally useful only when forecasts, historical context, hazard signals, uncertainty and recommended actions are connected without overstating what the data can prove. This updated review paper examines AERIS OPS as an implemented web prototype rather than as a validated disaster-forecasting service. The reviewed project combines a React/TypeScript client, Express/tRPC services, Drizzle persistence contracts, Open-Meteo forecast and air-quality adapters, ERA5-based historical analytics, deterministic multi-hazard screening, explainability and risk dashboards, Leaflet location intelligence, GDACS/NASA live hazard-layer adapters, agriculture guidance, alerts, incident review and PDF/CSV export. A key design feature is explicit provenance: outputs are labelled LIVE, FALLBACK, HISTORICAL, SAMPLE, RULE-BASED or UNAVAILABLE, while analytical triggers are kept separate from official warnings. The project snapshot passed TypeScript checking, production build, and 45 automated tests; these results demonstrate software-level verification, not predictive accuracy or real-world impact. The paper adds an implementation-specific research-gap matrix covering the integration-to-action gap, provenance and uncertainty gap, hazard data-availability gap, explainability evaluation gap, safe notification gap and prototype-to-evidence gap. Future work should establish labelled regional event datasets, calibrated models, exposure and vulnerability layers, human-centred evaluations, official operational integrations and prospective field validation.

***Index Terms***— weather intelligence, explainable AI, impact-based forecasting, disaster risk management, climate analytics, data provenance, rule-based screening, uncertainty communication, human-centred decision support.

## I. INTRODUCTION

Weather variables such as temperature, precipitation, humidity, wind, pressure, visibility and air quality are not themselves operational decisions. Their meaning depends on location, timing, exposure, infrastructure, vulnerability and the role of the user interpreting them. A forecast of intense rainfall, for example, can justify closer monitoring, but it cannot by itself establish flood depth, road closure, casualties or the need for evacuation. This distinction is central to safe weather-intelligence design.

AERIS OPS was reviewed against the uploaded project implementation and its accompanying engineering documentation. The current project is a browser-based command centre organised into weather, predictions, risk, explainability, decision support, disaster response, alerts, location intelligence, smart agriculture, air quality, trends, climate context, history and a disaster/climate lab. The implementation deliberately uses conservative language: it calls its hazard outputs screening indices and analytical advisories, not probabilities, official warnings or trained machine-learning predictions.

This paper has two purposes. First, it synthesises the research and operational principles relevant to the platform. Second, it updates the previous proposal into an implementation-aware review that distinguishes features present in the codebase, provider data actually connected, deterministic screens that are available, and evidence still required before operational deployment.

## II. PROBLEM IDENTIFICATION AND OBJECTIVES

The problem addressed by AERIS OPS is fragmentation across current weather, forecast series, historical climate context, air-quality values, hazard screening, official bulletins and response notes. Fragmentation increases the risk that users miss a change, confuse a forecast with an observation, interpret an analytical threshold as an official warning, or treat a missing dataset as evidence of low risk.

### *A. Objectives*

- Review explainable AI, impact-based forecasting, open weather services, climate analytics and human-centred operational interfaces.
- Map the actual implementation to a traceable data-to-action architecture.
- Identify research gaps that remain after the prototype has been implemented.
- Define validation measures for data quality, software correctness, predictive performance, explanation quality and user comprehension.
- Provide a responsible deployment boundary that preserves official-authority primacy and human oversight.

## III. LITERATURE AND OPERATIONAL FOUNDATIONS

### *A. Explainable AI for Risk and Disaster Management*

Explainable AI can expose the variables, rules or evidence that contribute to an output. In disaster-risk settings, explanations are useful only when they are faithful, stable, understandable and matched to the decisions users must make. AERIS OPS operationalises this principle through factor cards, evidence lists, source/time labels and separate explanation panels. Because the current hazard engine is deterministic, the explanation can show the threshold or missing-data condition directly. This is an important baseline before introducing opaque or high-capacity models.



### *B. Impact-Based Forecasting and Multi-Hazard Decision Support*

Impact-based forecasting shifts attention from atmospheric variables to possible consequences for people, assets and services. It requires exposure, vulnerability, infrastructure and local authority information in addition to weather. The project supports the transition only partially: it provides contextual screening, preparation guidance, incident review and decision-support workflows, but explicitly marks population exposure, road closures, shelters, hydrological conditions, terrain and verified affected boundaries as unavailable when those feeds are not connected.

### *C. Open Data, Reanalysis and Provider Integration*

The project uses Open-Meteo forecast and air-quality endpoints for current and forecast fields, and an Open-Meteo historical endpoint with ERA5 reanalysis for longer-range analytics. Reanalysis is model-assisted historical reconstruction and must not be described as a direct sensor observation. The implementation retains provider timestamps, source labels, coordinate validation, response checks, field coverage and stale-data logic. GDACS cyclone and flood event adapters and NASA GIBS VIIRS flood raster metadata extend the map layer while preserving the distinction between a point event and a verified flood extent.

### *D. Human-Centred Operational Interfaces*

Operational interfaces should reduce interpretation effort without hiding uncertainty or provenance. AERIS OPS follows an Atmospheric Signal visual language with a persistent command rail, evidence panels, severity bands, timestamp/provenance chips, maps and action-oriented guidance. The interface avoids relying on colour alone and states that empty or unavailable data is not equivalent to safe conditions.

## IV. REVIEW METHODOLOGY

This work uses a focused narrative review combined with implementation inspection. The original paper's themes were retained and cross-checked against the uploaded AERIS OPS project snapshot. Source inspection covered application routes, forecast and air-quality retrieval, disaster-prediction rules, historical signal analytics, trend aggregation, live hazard adapters, data schema, documentation and automated tests. The synthesis is not a systematic review, does not claim exhaustive database coverage, and does not perform a meta-analysis of forecasting accuracy.

Implementation verification was performed after installing the locked project dependencies. TypeScript checking completed successfully, the production Vite/Express build completed successfully, and the test suite reported 9 test files and 45 passing tests. These are software verification results. They do not establish calibration, generalisation, event-detection skill, operational availability or user benefit.

## V. AERIS OPS IMPLEMENTATION AND SYSTEM ARCHITECTURE

The implemented architecture is a layered web system. A single selected location is shared across routes, while server adapters and pure analytical helpers keep provider integration and screening logic testable. The architecture below is reconstructed from the project source and documentation.

| Layer | Implemented responsibility | Evidence / boundary |
|---|---|---|
| User and presentation | React command centre with weather, predictions, risk, explainability, alerts, GIS, agriculture, AQI, trends, history and response routes; Recharts and Leaflet. | Interactive dashboards and exports; browser UI is not emergency dispatch. |
| Application/ API | Express/tRPC routing, same-origin weather-trends, disaster-prediction and live-hazard adapters, validation and error handling. | Coordinates, horizon and payloads validated; upstream failures visible. |
| Data integration | Open-Meteo forecast and air quality; historical/ERA5; GDACS cyclone/ flood points; NASA GIBS VIIRS raster metadata; official-feed separation. | Source, retrieval time, coverage and limitations retained; no synthetic geometry. |
| Intelligence | Rule-based hazard screens, trend/ anomaly helpers, threshold alerts, evidence/factor explanations and highest-supported composite. | Scores are screening indices, not probabilities; trained ML unavailable until labelled data and a versioned model exist. |
| Persistence/ audit | Drizzle contracts for notification rules, alert events and simulation runs; local incident records and exports. | Outbound notification inactive without credentials, consent and an always-on worker. |

TABLE I. AERIS OPS IMPLEMENTATION LAYERS AND BOUNDARIES

### *A. Data Provenance and State Semantics*

A central contribution of the implementation is treating data quality as an output. LIVE denotes a recent provider response that passed validation; FORECAST denotes future provider values; HISTORICAL denotes archive or reanalysis values; FALLBACK/ SAMPLE denotes illustrative or cached content; RULE-BASED denotes deterministic screening; and UNAVAILABLE indicates that required coverage is missing or a provider is not usable. This vocabulary reduces the risk of silently replacing missing values with plausible numbers.

### *B. Hazard Screening and Missing-Data Handling*

The multi-hazard route contains ten hazard modules and uses explicit evidence and data-gap descriptions. Supported weather-context screens include rainfall, heat, thunderstorm, fog/visibility and selected multi-hazard conditions. Drought, cyclone landfall, landslide susceptibility, official flood extents, district-wide exposure and other specialist outputs remain unavailable when required hydrology, terrain, track, event, warning-geometry or exposure data is absent. The composite uses the highest supported screen and counts additional elevated signals rather than averaging incomparable hazard scores.

### *C. Historical Analytics and Live Hazard Layers*

Historical intelligence aggregates daily weather-derived proxy screens such as high temperature, heavy rain, strong gusts, severe weather codes and low visibility when observations exist. Trends and anomalies use explicit sample-count and coverage requirements; nulls are retained and no district-wide or India-wide surface is inferred from one selected point. The GIS layer represents GDACS flood results as event points, GDACS cyclone observations as time-stamped points and NASA VIIRS products as satellite raster coverage. The implementation does not draw invented circles, polygons or tracks from isolated observations.

### *D. Domain Modules and Operational Workflow*

Smart Agriculture translates live weather into transparent farm-care guidance for crop, stage, soil, irrigation and field context; it does not claim yield gains or prescribe pesticides. AQI Intelligence shows the selected Open-Meteo US AQI point and particulate values while stating that US AQI is not CPCB AQI and that verified historical station coverage is unavailable. Disaster Response records local incidents and recovery tasks for review, but does not claim dispatch, shelters, hospitals, road closures, casualties or verified damage.



## VI. RESEARCH GAP

The previous version identified a general distance between data access and actionable information. The implementation review makes that gap more precise. AERIS OPS addresses several design aspects, but the following research gaps remain open and should guide the next research phase.

### *A. Theoretical Basis of the Research Gap*

The research gap can be understood through a socio-technical data-to-action chain: observation, data quality control, contextual interpretation, risk communication, human decision, and monitored action. A failure at any stage can reduce safety even when the underlying forecast is technically accurate. For example, a timely rainfall forecast may still produce a poor decision if its timestamp is unclear, its spatial coverage is misunderstood, its uncertainty is hidden, or the user cannot connect it to local exposure and available response options.

A second theoretical issue is the distinction between hazard, exposure, vulnerability and impact. Hazard describes a potentially damaging physical process; exposure describes people, assets or services located in the affected area; vulnerability describes susceptibility to harm; and impact is the realised consequence. A weather API primarily supplies hazard-related variables. Converting those variables into an impact statement requires additional geospatial, demographic, infrastructure and event data. Therefore, a point-weather threshold cannot be treated as a complete disaster prediction.

Uncertainty also has more than one source. Aleatoric uncertainty arises from natural atmospheric variability, while epistemic uncertainty arises from incomplete observations, model limitations, missing exposure data or uncertain provider coverage. A safe interface should communicate both types through forecast horizons, freshness, confidence limitations, missing-data states and source labels. The AERIS OPS use of LIVE, FORECAST, HISTORICAL, SAMPLE, RULE-BASED and UNAVAILABLE states provides a design baseline, but the effect of these labels on user interpretation still requires empirical study.

Explainability provides a further theoretical boundary. An explanation may describe why a rule fired, but it does not prove that the rule is causally correct or that the predicted event will occur. Meaningful evaluation must therefore examine fidelity to the implemented decision logic, stability under small input changes, completeness of important factors, usefulness for counterfactual reasoning and whether users develop calibrated—not excessive—trust. This distinction separates explanation quality from model accuracy and from visual clarity.

| Gap | Existing work often provides | AERIS OPS addresses | What remains |
|---|---|---|---|
| G1. Integration-to-action | Forecast variables, maps or alerts in separate tools. | Connects weather, historical context, risk screens, explanations, guidance, incident review and export. | Measure task completion, decision speed and appropriate escalation against existing practice. |
| G2. Provenance and uncertainty | Values without clear source, age, transformation or missing-data semantics. | LIVE/FORECAST/ HISTORICAL/ SAMPLE/ UNAVAILABLE and provider disclosures. | Test whether users correctly interpret freshness, reanalysis and uncertainty. |
| G3. Hazard data availability | Implied complete coverage despite missing exposure, terrain, hydrology or official geometries. | Unavailable states; no fabricated footprints, tracks, population impacts or warnings. | Acquire regional datasets and quantify coverage-dependent failure modes. |
| G4. Explainability evaluation | Explanations judged by visual appeal or trust alone. | Deterministic factors, thresholds, evidence and data gaps. | Evaluate fidelity, stability, completeness, counterfactual usefulness and comprehension. |
| G5. Safe notification | Analytical triggers blurred with authoritative warnings; duplicate messages. | Separates official bulletins; documents consent, secrets, deduplication and disabled-by-default delivery. | Validate signed sources, rate limits, receipts and opt-out handling. |
| G6. Prototype-to-evidence | Working dashboard mistaken for validated predictive model. | 45 passing tests, successful type-check/ build and explicit non-claims. | Conduct chronological event evaluation, calibration, prospective trials and impact evaluation. |

TABLE II. IMPLEMENTATION-SPECIFIC RESEARCH GAPS

### *A. Main Research Gap and Proposed Research Question*

The main research gap is the lack of a validated, locally contextualised and human-evaluated bridge between heterogeneous weather data and safe operational action. Existing systems tend to optimise one part of the chain—forecasting, mapping, historical analysis, explainability or official warning dissemination—while leaving integration, provenance, data gaps and user interpretation to the operator. AERIS OPS provides a concrete prototype for studying this bridge, but it does not yet provide the empirical evidence needed to claim that the bridge is accurate, trusted or beneficial.

**Proposed research question:** How can a provenance-aware, uncertainty-explicit and explainable multi-source weather-intelligence interface improve the accuracy, timeliness and safety of location-specific preparedness decisions, while reducing false confidence when hazard or exposure data is unavailable?

## VII. EVALUATION FRAMEWORK

Evaluation should be separated into five layers so that a passing software test is not confused with a successful forecast or useful operational intervention. The framework follows a layered validity model: construct validity asks whether the screen represents the intended hazard concept; internal validity asks whether the implementation produces the intended output from the available inputs; external validity asks whether results transfer across locations, seasons and provider conditions; and operational validity asks whether users can make safer decisions with the system.

### *A. Software and Data-Quality Evaluation*

Software evaluation should test valid and invalid coordinates, forecast horizons, malformed provider payloads, missing fields, stale timestamps, timezone conversion, unit handling, null preservation, API timeouts and fallback behaviour. Regression tests should confirm that analytical alerts are not generated from stale or unavailable observations, that official bulletins remain separate from



app-generated triggers, and that a missing value is not silently converted to zero. Data-quality monitoring should report completeness, freshness, range validity, spatial coverage, temporal alignment, provider response rate and the proportion of outputs labelled UNAVAILABLE.

### *B. Predictive and Statistical Evaluation*

If a future trained model is added, event labels and forecast horizons must be defined before modelling. Discrimination measures such as precision, recall, F1 and precision-recall area under the curve should be reported together with calibration measures such as the Brier score and reliability diagrams. Lead time, false-alarm rate, missed-event rate and performance by hazard class are important because a model can obtain acceptable average accuracy while failing on rare but high-consequence events. Chronological splits, geographically separated test sets and comparison with persistence, climatology and threshold baselines are necessary to reduce leakage and inflated performance.

### *C. Explainability, Usability and Human Factors*

Explainability evaluation should measure whether users can identify the main contributing factors, distinguish forecast from observation, recognise an unavailable dataset, and select an appropriate next action. Usability studies can measure task completion time, error rate, uncertainty comprehension, correct escalation, unnecessary escalation, information-retrieval accuracy and trust calibration. Participants should represent the intended roles, such as emergency managers, municipal staff, farmers, researchers and general users. Accessibility, language, colour vision, mobile display and low-connectivity conditions should be included in the test plan.

### *D. Operational and Ethical Evaluation*

Operational evaluation should use shadow-mode deployment before live notification. The system can generate internal analytical records while trained operators compare them with official bulletins and observed events. Evaluation should record whether a message was timely, understandable, actionable and appropriately qualified. Ethical review should address privacy in notification records, consent for SMS or email, auditability, unequal coverage, automation bias and the risk that users treat a screening score as an evacuation order. These safeguards are essential because a technically correct interface can still create harm through overconfidence or inappropriate escalation.

| Layer | Example measures | Current status |
|---|---|---|
| Software correctness | Type-check, build, unit/API contract tests, invalid-input and provider-failure tests. | Type-check/build passed; 45/45 tests passed. |
| Data quality | Completeness, range checks, freshness, timestamp alignment, coordinate validity, source attribution and null preservation. | Implemented in adapters and analytics; regional monitoring remains. |
| Predictive performance | Precision, recall, F1, PR-AUC, Brier score, calibration, lead time and false-alarm rate. | Not reported; no validated trained-ML model or labelled regional benchmark. |
| Explanation/ usability | Fidelity, stability, task completion, uncertainty comprehension, escalation and trust calibration. | Design intent present; controlled user evaluation is future work. |
| Operational impact | Missed warnings, unnecessary escalation, response time, equity/ accessibility and outcomes. | Not measured; requires approved field study and partners. |

TABLE III. EVALUATION LAYERS AND CURRENT EVIDENCE

### *A. Recommended Experimental Protocol*

- Construct a regional event catalogue with clear definitions for flood, heat, thunderstorm, fog and other target hazards.
- Pre-register forecast horizons, label definitions, exclusion rules and missing-data treatment.
- Split data chronologically and geographically where possible to reduce leakage and test transferability.
- Compare persistence, climatology, threshold and statistical baselines before machine-learning models.
- Report calibration and false-alarm costs in addition to discrimination metrics.
- Run explanation tests with meteorologists, emergency managers, farmers and general users.
- Conduct prospective shadow-mode trials before any alert or response automation is enabled.

## VIII. DISCUSSION

The reviewed implementation demonstrates that responsible weather intelligence is as much a data-governance problem as a modelling problem. AERIS OPS makes conservative choices that are academically valuable: it does not label illustrative fallback values as live; it keeps forecast context separate from official bulletins; it treats historical reanalysis as historical; it refuses to infer flood extent from rainfall alone; and it exposes data gaps instead of converting them to reassuring low-risk values.

The principal trade-off is that conservative availability states can make a dashboard look less complete than a system that fills gaps with assumptions. For safety-critical decision support, this is desirable. However, “UNAVAILABLE” must lead to an actionable next step—consulting an official source, obtaining a missing layer, or recording an observation—rather than becoming a dead end. The next research phase should study both the accuracy of signals and the quality of human decisions made around them.

## IX. LIMITATIONS AND FUTURE SCOPE

The project depends on third-party provider availability, gridded forecast uncertainty, uneven regional coverage and the quality of external official feeds. Current screens do not establish event probability, flood depth, cyclone landfall, landslide susceptibility, drought severity, exposure, casualties or economic loss. The project snapshot also includes fallback/sample states and simulation-oriented structures; these are useful for interface development but must remain clearly labelled in deployment.

Future work should add verified regional observation and event datasets, hydrology and terrain layers, exposure/vulnerability information, multilingual and accessible explanations, model versioning, calibration monitoring, official warning integration, authenticated notification workers, role-based audit trails and prospective user studies. A trained model should be introduced only when a labelled dataset, reproducible training process, versioned artifact, chronological evaluation and documented failure analysis are available.

## X. ADVANTAGES AND APPLICATIONS

- Unifies weather, historical context, risk screening, explanations and decision-support records.
- Makes source, freshness and data gaps visible at the point of use.
- Supports weather-aware agriculture and environmental exposure review without claiming yield or health outcomes.
- Provides a modular testable baseline for future statistical or machine-learning models.
- Creates an audit-friendly separation between analytical screens, official bulletins and local review records.

Potential applications include education, community preparedness, municipal situation monitoring, logistics, event planning, agricultural timing and research prototyping. Emergency operations use should require validated official data, approved procedures, trained users, accessibility review and human authorisation.

## XI. CONCLUSION

AERIS OPS is best understood as a provenance-aware weather-intelligence and decision-support prototype, not as a completed disaster-prediction service. The updated review shows that its



strongest contribution is the explicit connection of provider data, historical context, rule-based evidence, explanations, action guidance and data-status semantics in one interface. The codebase provides meaningful software evidence—successful type-checking, production build and 45 passing automated tests—but does not yet provide empirical evidence of predictive skill or operational impact. The central research opportunity is to evaluate whether transparent, uncertainty-aware integration can help people act earlier and more safely without creating false confidence when essential hazard and exposure data is missing.

## REFERENCES

[1] "Explainable artificial intelligence in disaster risk management: Achievements and prospective futures," *International Journal of Disaster Risk Reduction*, vol. 98, art. 104123, 2023.

[2] World Meteorological Organization, "Impact-based Forecast and Warning Services."

[3] Open-Meteo, "Weather Forecast API Documentation."

[4] Open-Meteo, "Historical Weather API Documentation."

[5] Open-Meteo, "Air Quality API Documentation."

[6] United Nations, "Early Warnings for All."

[7] United Nations Office for Disaster Risk Reduction, *Sendai Framework for Disaster Risk Reduction 2015–2030*, 2015.

[8] World Meteorological Organization, *Multi-hazard Early Warning Systems: A Checklist*, 2018.

[9] Intergovernmental Panel on Climate Change, *Climate Change 2022: Impacts, Adaptation and Vulnerability*, Cambridge University Press, 2022.

[10] H. Hersbach *et al.*, "The ERA5 global reanalysis," *Quarterly Journal of the Royal Meteorological Society*, vol. 146, no. 730, pp. 1999–2049, 2020.

[11] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you?: Explaining the predictions of any classifier," *Proc. ACM SIGKDD*, 2016, pp. 1135–1144.

[12] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," *NeurIPS* 30, 2017.

[13] F. Doshi-Velez and B. Kim, "Towards a rigorous science of interpretable machine learning," arXiv:1702.08608, 2017.

[14] Food and Agriculture Organization, "Climate-Smart Agriculture."

[15] D. S. Wilks, *Statistical Methods in the Atmospheric Sciences*, 4th ed., Elsevier, 2019.

[16] T. Gneiting and A. E. Raftery, "Strictly proper scoring rules, prediction, and estimation," *JASA*, vol. 102, no. 477, pp. 359–378, 2007.

[17] F. Pappenberger *et al.*, related operational machine-learning forecasting literature, used as a methodological comparator for flood/event evaluation.

[18] M. Reichstein *et al.*, "Deep learning and process understanding for data-driven Earth system science," *Nature*, vol. 566, pp. 195–204, 2019.

[19] AERIS OPS project snapshot, README.md, package.json, source tree and automated tests, uploaded ZIP archive, 2026.

[20] AERIS OPS project documentation, "Integration source research," *docs/integration-sources.md*, uploaded ZIP archive, 2026.

[21] AERIS OPS project documentation, "Alert notification setup," *docs/alert-notifications-setup.md*, uploaded ZIP archive, 2026.

[22] AERIS OPS project plan, *plan.md*, implementation decisions and evidence boundaries, uploaded ZIP archive, 2026.

## APPENDIX A. IMPLEMENTATION TRACEABILITY SUMMARY

| Paper claim | Project evidence inspected | Responsible interpretation |
|---|---|---|
| Location-based weather intelligence | App.tsx; WeatherAnalyticsPanel.tsx; Open-Meteo forecast calls. | Provider-backed current/ hourly/daily context; not a sensor network. |
| Rule-based hazard screening | shared/disaster-predictions.ts; server/ disaster-predictions.ts; prediction tests. | Deterministic screening indices; not probabilities or trained ML. |
| Historical analytics | server/weather-trends.ts; shared/ weather-trends.ts; historical signal tests. | ERA5/reanalysis and proxy screens; not verified disaster-event history. |
| Live hazard map context | server/live-hazards.ts; live-hazards.test.ts; integration-sources.md. | GDACS points and NASA raster metadata; no fabricated geometry. |
| Explainability and safe workflow | ExplainableAIDashboard.tsx; RiskIntelligenceDashboard.tsx; alert docs. | Evidence and provenance visible; user studies and live delivery remain future work. |
| Engineering verification | 9 Vitest files; 45 passing tests; type-check and production build. | Software-level evidence only; no claim of predictive or operational validation. |

TABLE IV. IMPLEMENTATION TRACEABILITY SUMMARY