databricks

RESPONSE TO MARKET CONSULTATION

AIWaM02 Forecast Assistant — Market Consultation

Prepared for Rijkswaterstaat (RWS)

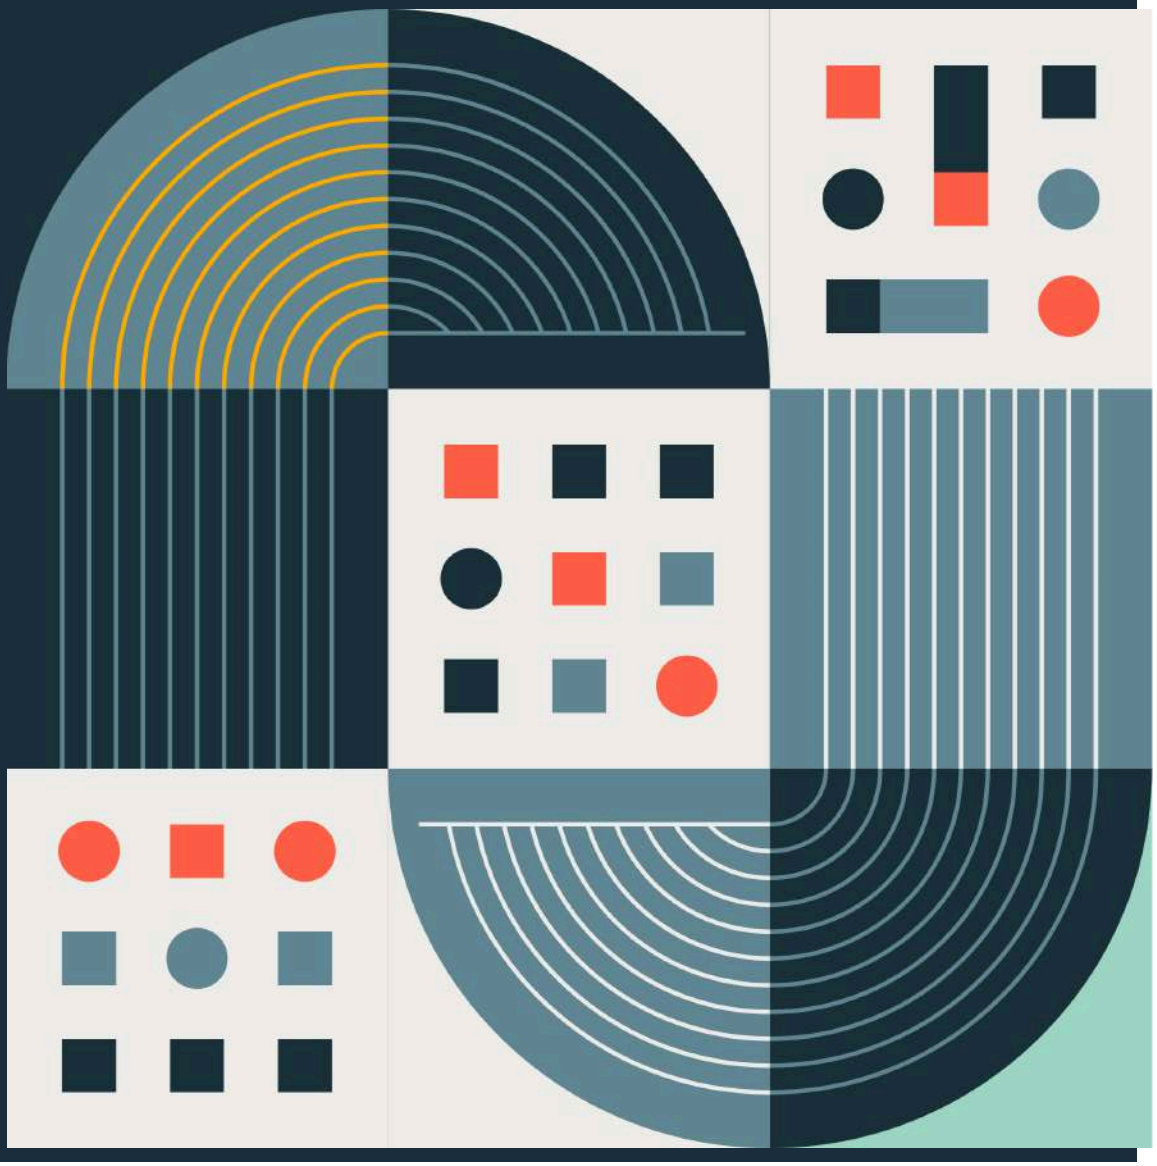

Confidential

Organisation and contact details

Response to the market consultation for the AIWaM02 Forecast Assistant of Rijkswaterstaat (RWS).

Item	**Details**
Organisation	Databricks Nederland B.V.
Contact person	Jasvant Gadjiedas, Enterprise Account Executive Public Sector
Telephone	+31 6 48 55 55 20
Email	jasvant.gadjiedas@databricks.com
Language	Dutch / English

Reading guide

Databricks is a data and AI platform software company, and we answer this consultation from that position. The Forecast Assistant would be a **configured solution on our generally available Data + AI platform**, delivered together with a hydrological domain partner that provides operational forecasting expertise. Where a question concerns delivery detail or hydrological specifics, we answer in terms of how the platform and model layer support them, and leave the delivery detail to the joint team.

Vision & Solution Overview

Our vision is to empower Rijkswaterstaat forecasters with a modern, transparent Data + AI Forecast Assistant that integrates diverse hydrological datasets, meteorological inputs, and expert operational knowledge into actionable lines of reasoning ('redeneerlijnen'). Built on Databricks' robust data and AI platform, our proposed solution preserves full human oversight while streamlining analysis, enhancing uncertainty communication, and enabling seamless decision support across all forecasting phases.

One-page summary

AIWaM02 Forecast Assistant — Market Consultation · Prepared for Rijkswaterstaat (RWS)

What we propose. An AI Forecast Assistant that *supports* forecasters rather than replacing them: it integrates operational model output, observations, meteorology, upstream notifications and historical reports into a single, traceable line of reasoning (*redeneerlijn*). It explains and contextualises existing forecasts — it never produces a competing forecast value and is never in the critical path of the operational chain (FEWS). It is a configured solution on our generally available, governed Data + AI platform (Unity Catalog, Lakeflow, AI Search, Agent Bricks + MLflow, model gateway, Databricks Apps), not a black-box product.

Why it fits

- **Expert-in-the-loop by design** — every statement labelled by origin (data / calculation / expert rule / AI) and cited with source and issue time; uncertainty always shown; human approval before anything leaves the tool.
- **Trustworthy under extremes** — numbers come only from tested tools, out-of-range states are flagged, and validation uses honest hindcast replay with strict "as-of" time control.
- **Sovereign & compliant** — EU data residency, RWS-owned storage and keys, private networking, BIO/ISO 27001, GDPR, EU AI Act readiness; no RWS data used to train third-party models.
- **Open & low lock-in** — open formats (Delta/Iceberg), MLflow, MCP and a model-agnostic gateway; data, prompts, evaluation sets and knowledge stay under RWS control with a clear exit path.

Recommended approach

- **Procurement:** innovation partnership, one integrated team, PoC → operation without a new tender; don't split into lots.
- **Phasing (≈3–5 yr):** PoC (4–6 mo) → Pilot in shadow mode through a high-water season (6–12 mo) → Operation (2–4 yr), with go/no-go gates.
- **Indicative platform consumption** (draft, not validated): PoC €2–5k/mo · Pilot €5–15k/mo · Operation €8–25k/mo (build & maintenance priced separately).

What we need from RWS. Read access to RWsOS/Delft-FEWS forecasts, ensembles and the archive of *issued* forecasts (MATROOS); observations; report archives; ~0.5 day/week each from forecasters and senior hydrologists; a cloud environment under RWS control with security approvals; and clear decision moments with a product owner.

Key differentiator. A single governed platform that already provides the building blocks, so effort goes into hydrological reasoning, knowledge capture and evaluation — not into assembling tools. Evaluation-first delivery with expert-approved historical cases makes trustworthiness measurable before operational use.