

# An Interoperable AI driven Digital Twin Ecosystem for Water Resource Management and Anticipatory Action in the Global South. A Case Study for Limpopo River Basin

Mariangel Garcia Andarcia , Hugo Retief \*, Davy Vanham, Moctar Dembele, Keerththanan Vickneswaran, Emmanuel Ayodele Jolaiya, Zolo Kiala, Abdul Afham, Paulo Silva, Chris Dickens, Pedro Chambel-Leitão, Flávio Santos, Débora Barreiros, Hélio Santos, Karthikeyan Matheswaran, Lal P Muthuwatta, Nicholas Pattison, Kayathri Vigneswaran, Henry Roman, Samuel Storr, Surajit Ghosh

*International Water Management Institute (IWMI), Colombo, Sri Lanka*

*\*Corresponding author: h.retief@cgiar.org*

## ABSTRACT

Water management in the Global South increasingly requires integrated digital systems that can support both strategic planning and anticipatory action across transboundary basins. Fragmented dashboards, disconnected models and static reports remain insufficient where hydrometeorological data are sparse, institutions span multiple jurisdictions, and decision windows are short. This paper presents the architectural design and early operational implementation of a modular, service-oriented digital twin ecosystem developed by IWMI and demonstrated in the Limpopo River Basin in Southern Africa, in co-development with the Limpopo Watercourse Commission (LIMCOM).

The ecosystem couples five interoperable layers through open application programming interfaces (APIs): a data infrastructure based on an Open Data Cube and SpatioTemporal Asset Catalog (STAC) organising Earth observation, hydrometeorological and ecological datasets; operational hydrological services including SWAT+ outputs, reservoir surface-area monitoring, drought analytics and flood-exposure products; environmental-flow risk and ecological monitoring components such as PROBFLO-aligned assessments and MiniSASS citizen-science observations; a basin digital twin geospatial interface using TerriaJS-based cataloguing; and WaterCopilot, a natural-language assistant that uses retrieval-augmented generation (RAG) and validated tool-calling against the same APIs serving the geospatial layer, ensuring a single source of analytical truth.

Early operational evidence indicates that the architecture supports daily-to-seasonal decision workflows for transboundary stakeholders, while exposing characteristic limitations: hydrological model calibration uncertainty in data-scarce subbasins, Earth observation latency, variable quality of citizen-generated observations, and the need for continuous validation of AI-mediated responses. We argue that the principal scientific contribution is not any single model, dashboard or assistant, but the integration of heterogeneous hydrological information into a service-oriented architecture whose components can be independently improved, replaced or extended. The approach offers a transferable framework for operational water resource management and anticipatory action in data-constrained Global South basins, conditional on open standards, institutional co-development, modular APIs and sustained validation.

**Keywords:** Hydroinformatics; Service-oriented architecture; Retrieval-augmented generation; WaterCopilot; Citizen science; SWAT+; Transboundary water governance.