26th International Congress on Modelling and Simulation, Adelaide, SA, Australia, 30 Nov–4 Dec 2025
mssanz.org.au/modsim2025

FIRM: cloud-native integration framework for basin-scale hydrological modelling in the Murray-Darling Basin

P Gijsbers[1], S De Kleermaeker[1], C Urich[2], J Rahman[3], J Fuller[4], E Brennan[4], L Peralta[4], D Vorreiter[4], B Cressall[2], N Price[5], S Dixon[5]

[1] Deltares, [2] HARC, [3] FlowMatters, [4] Murray Darling Basin Authority, [5] A23
Email: Simone.DeKleermaeker@Deltares.nl

Abstract: This paper introduces FIRM, a cloud-native framework designed to support basin-wide integrated hydrological modelling in the Murray-Darling Basin. The Basin is a vital national resource in Australia, but its management is complex due to the fragmented governance structure involving multiple states and basin jurisdictions. Currently, hydrological models are developed and maintained independently by each jurisdiction, making basin-wide analysis time-consuming and difficult to reproduce. To address these challenges, the Murray Darling Basin Authority (MDBA) has developed a Framework for Integrated River Modelling (FIRM), which enables the integration of diverse hydrological models, facilitates collaboration between jurisdictions, and supports advanced scenario analysis. FIRM is built on several technological components: Delft-FEWS (a workflow engine for model orchestration), FOM (a python based tool for automated model integration), and FRAME (a web-based dashboard for visualisation). These components operate within MDIRAS (Murray-Darling Integrated River Assessment System), a Microsoft Azure environment, providing dedicated workspaces for each jurisdiction as well as a collaboration hub.

A key feature of FIRM is its metadata model, which ensures transparency and reproducibility. Models are treated as black boxes maintained by the Jurisdictions to preserve their integrity, and metadata is used to describe model inputs, outputs, and interconnections. The metadata model includes model details, timeseries mappings, and linkage definitions. It enables the support of flexible scenario configurations across four dimensions: basin scope, reference scenarios (e.g., Without Development, Basin Plan 2024), connectivity between models (e.g., disconnected, feedback loops), and functionality (e.g., climate scaling, experiments). Users can select and run (a combination of) models and analyse results, with full traceability of inputs, outputs, and user actions.

To date 88 models across 4 reference scenarios have been integrated into FIRM, using 26 model code versions. The models can be run alone or linked together, for historical conditions as well as climate change scenarios. The implementation of FIRM faced several integration challenges, including inconsistent standards across hydrological software, technical limitations in command-line interfaces, and changes in data formats between model versions. As it stands, the framework represents a significant step forward in enabling collaborative, transparent, and reproducible hydrological modelling across the Basin.

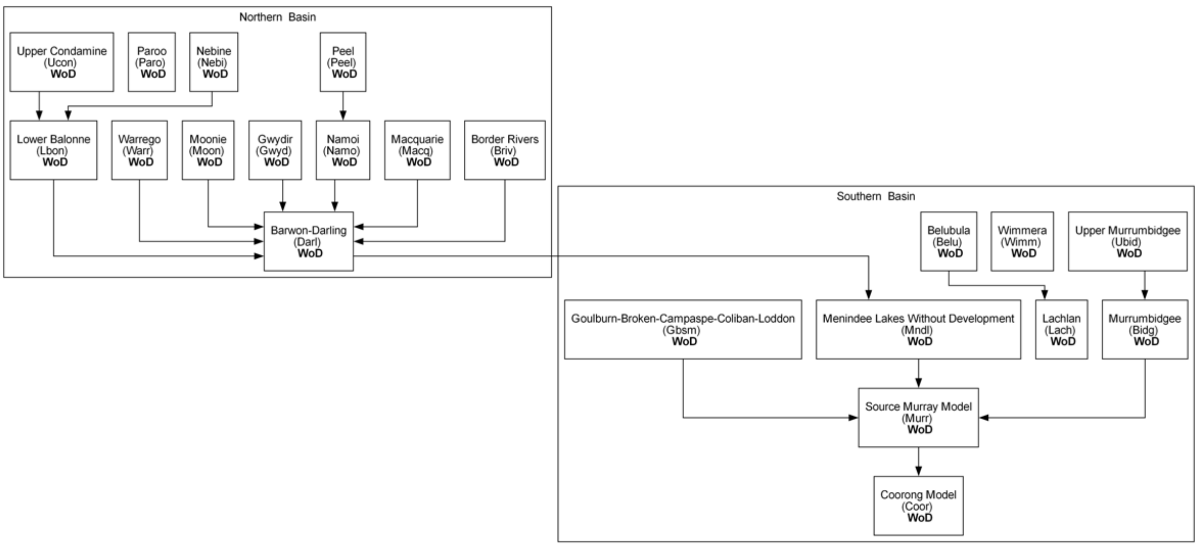

Figure 1 Model linkage diagram for Without Development (WoD) Reference Scenario, as generated by 'FOM'

Keywords: *Hydrology, Model integration, Cloud-native, Scenario Analysis, Murray-Darling*

1. INTRODUCTION AND MOTIVATION

The Murray–Darling Basin (MDB) is a national resource, valued by all Australians. Towns and rural communities across the Basin rely on a healthy river system. The importance of the Basin is matched by the challenge of its management. Managing the health of the river system and balancing the competing demands of the water in the system is an important and complex task. This requires coordination and cooperation between the Commonwealth, Basin Jurisdictions, local communities, industry, and other stakeholders with varying interests in Basin outcomes.

1.1. Many states, one basin: a complex governance challenge

Management of water in the MDB is primarily the responsibilities of Basin jurisdictions who maintain their own models of the Basin catchments, with the Murray-Darling Basin Authority (MDBA) given responsibility by the states to maintain the model of the Murray and Lower-Darling. Availability of integration frameworks to support multi-catchment modelling requirements is limited, especially for this situation where **legacy models** and **many software versions** need to be combined. Despite its importance, current systems and tools supporting water resource modelling of Basin hydrology are not keeping up with the level of analysis required to understand the increasingly complex factors that are impacting the state of the Basin now, and into the future. Models are not integrated across the Basin, so a holistic analysis can take months of manual effort to produce meaningful, timely, and comprehensive reporting. Moreover, the individual river system models are not yet fully integrated across the Basin. This disaggregated approach results in several limitations hindering basin wide studies.

1.2. Objectives

A statutory requirement under section 127 of the Water Act (2007) is for the MDBA 'to develop, in consultation with the Basin States, an integrated water model for the Murray-Darling Basin.' To achieve this and to undertake Basin wide studies, the MDBA and the States need a **reproducible, shared platform** for collaboration on hydrological models. The Integrated River System Modelling Framework (IRSMF) was developed by the CSIRO for the Murray-Darling Basin Sustainable Yields project in 2007/08 [Yang 2010], was adopted by the MDBA with improvements to develop the Original Basin Plan. To achieve the present-day targets of Basin Plan review, there was a need for a new platform to represent the current workflows and models, which has been achieved through the Integrated River Modelling Uplift Program [MDBA 2022].

The updated platform should provide a reproducible system with a **robust workflow** to connect and integrate the basin catchment models and to ingest selected data sets. The platform has to enable **secure sharing** of models and outputs, and support the MDBA in its efforts to mature, integrate, and uplift its and the Basin jurisdictions' current data modelling, data management and analytics functionalities, and capabilities. This paper introduces **FIRM**, the <u>F</u>ramework for <u>I</u>ntegrated <u>R</u>iver <u>M</u>odelling, which is developed to facilitate this transparent and reproducible workflow for model integration and execution.

1.3. Use case: Basin Plan Review

FIRM will be used to support the Basin Plan Review process, which requires assessment of different Reference Scenarios (i.e. development scenarios), each involving different valley models and linkages. Figure 1 shows the Without Development Reference Scenario (WoD), where all the dams, irrigation and environmental infrastructure, all consumptive users as well as any management and operational rules and actions are removed or disabled from the river system models. This Reference Scenario only requires direct linkages from an upstream model to a downstream model, while other Reference Scenarios, such as Basin Plan 2024, include management and operational rules or irrigation diversion which generate feedback loops (i.e. information flowing upstream) that need to be iterated through as the workflow is executed with the user able to select how many iterations are run in a particular workflow.

The FIRM metadata model (section 2.2), defining the Reference Scenario and valley Linkages, allow different specific configurations to be tested. Diagrams like the one in Figure 1 are automatically generated for each Reference Scenario based on the metadata as part of the model onboarding process with FOM (section 2.3).

2. FIRM – ENVIRONMENT, TECHNOLOGY AND APPLICATION

2.1. Environment to collaborate

To facilitate the collaboration on hydrological models in the Murray-Darling Basin, MDBA has established an Azure environment called **MDIRAS** (<u>M</u>urray-<u>D</u>arling <u>I</u>ntegrated <u>R</u>iver <u>A</u>ssessment <u>S</u>ystem). Figure 2 shows the different workspaces within MDIRAS. Each jurisdiction is offered its own <u>Jurisdiction Spoke</u>, i.e. an Azure DevOps project dedicated Git repositories for model development and a private computing environment for development with virtual workstations and storage. Models intended for collaboration across the basin are released to the <u>Collab Hub</u> for collaboration. For models selected for integration, meta data files are added, automated integration test runs executed, and (on success) the model is added to the <u>FIRM Hub</u> environment. How these environments are used is described in the MODSIM 2025 paper "Shifting Left: Building a collaborative integrated hydrological model for the Murray-Darling Basin" [Brennan et al. 2025].

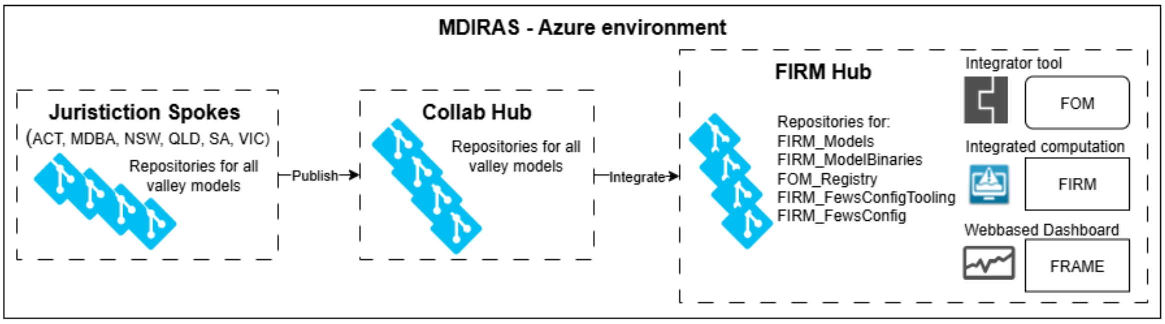

Figure 2 Major workspaces in the MDIRAS environment related to FIRM

Once a model is available in the FIRM Hub, it can be onboarded using **FOM,** the FIRM Onboarding Manager**,** a tool which automatically integrates the model for use within the FIRM application. The **FIRM metadata model** enables this automated generation of the FIRM application. The **Model integration workflows** facilitate the integration of the hydrological models into one combined fully traceable and reproducible Scenario, as defined, run and analysed in Delft-FEWS, the software underlying the **FIRM application**. Results can be disseminated through **FRAME**, a web-based dashboard. All these components are described in more detail below.

2.2. FIRM metadata model

To enable transparent and reproducible integration of models across jurisdictions, FIRM employs a structured metadata model of the Murray–Darling Basin. This model provides a consistent description of each contributing model and its relationships, while preserving the integrity of individual models and the autonomy of the organisations that maintain them. By making model linkages explicit, the metadata framework builds trust in integrated Basin simulations and their outcomes.

In FIRM, each model is assigned to a valley (i.e. catchment) and a reference scenario. Valleys are conceptual units that consume and generate time series data. The models within them are treated as black boxes, which basically remain untouched except for small updates to path references and use of reload functionality. Time series are associated with a specific site and parameters are drawn from a global list, ensuring a consistent mapping across different model outputs. Linkages, as illustrated in Figure 3, describe how valleys, sites, and parameters connect, whether unidirectionally (e.g. downstream flows) or bidirectionally (e.g. upstream water orders/allocations), with feedback loops resolved through iterative runs. Sites correspond one-to-one to FIRM locations in the Delft-FEWS administration, except for sites that act as model linkage. Linkage sites are represented by two FIRM locations, one location for the (data) providing valley and one location for the (data) receiving valley. This separation allows a verifiable data transfer.

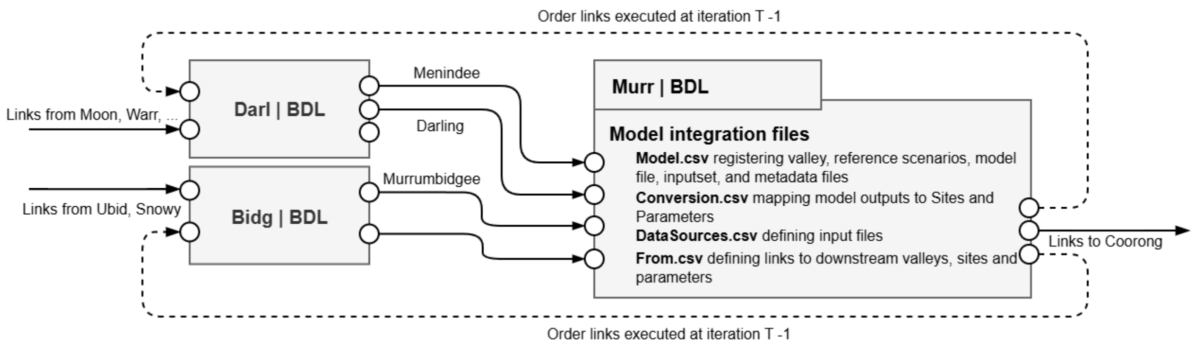

Figure 3 Linkage meta data: example of bi-directional linkage between Darling, Murrumbidgee and Murray model

The metadata is defined at the level of the individual valley (i.e. model), requiring only limited knowledge of upstream or downstream systems. This significantly reduces integration complexity, increases transparency, preserves each jurisdiction's autonomy to maintain and integrate its own models, and allows for secure sharing of models if required.

Metadata components

The metadata model comprises three main components at the individual model level:

- **Model metadata** (Model.csv) – Valley, reference scenario, model code and version, Git repository hash, dataset location, and model data files.
- **Time series metadata** (Conversion.csv, DataSource.csv) – Id mapping between model in- and outputs and FIRM administration (sites/locations and measurands/parameters); identification of inputs to be scaled in climate scenario runs.
- **Linkage metadata** (From.csv) – Data transfer definition between outputs from this (providing) valley and the inputs of receiving valleys, for both downstream flows and upstream orders.

All metadata is stored in a FIRM integration folder as CSV files together with the model, forming the backbone of FIRM's integration process and providing change transparency.

2.3. FOM – Integration Tool

FOM (FIRM Onboarding Manager) is a Python-based tool that operationalises the FIRM metadata to assemble the integrated Basin model in the FIRM application (see Figure 4). At the model scale, FOM interrogates the supplied files to generate integration files, verify that models can be run successfully, confirms that outputs are generated, and summarises these checks in a QA report. At the Basin scale, it draws on the model registry to compile the full Delft-FEWS configuration of FIRM, ensures all linkages between models are valid, tests that the system will run in the cloud environment, and produces a basin and scenario wide QA summary and documentation.

In short, FOM ensures that when a user creates a scenario in the FIRM application, the selected model version for each valley will apply the mappings and linkages suitable for the chosen reference scenario, as defined by the metadata model.

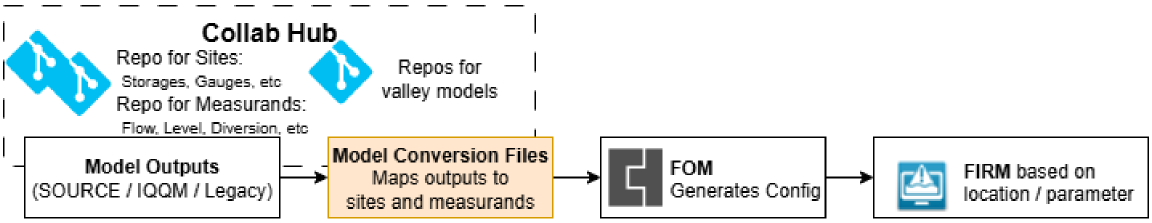

Figure 4 Timeseries meta data, mapping model results of all models to the correct location and parameter

FOM can operate as a standalone utility for local setups, or as part of a CI/CD workflow. In Continuous Integration, it provides early feedback on integration issues, automates the compilation of the Basin model, and updates the operational FIRM configuration (see Figure 5). By combining metadata from the model registry with user-specified integration details, FOM produces the configuration that governs workflows, data exchange, and time series administration across the Basin.

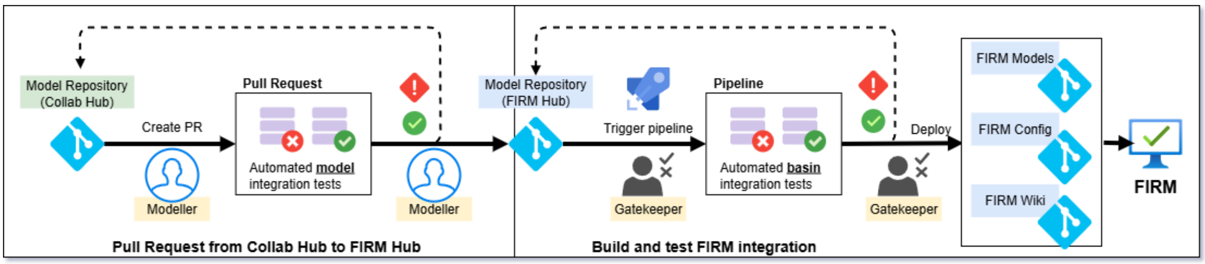

Figure 5 Integration of a model from Collab Hub into FIRM Hub, steps and processes

FOM currently supports the integration of models developed in 6 software packages: Source, IQQM, REALM, and the legacy St George, Snowy, and Coorong models.

2.4. Model integration workflows

Models in FIRM application can be run individually or in combination with other models in the same reference scenario. Models can be run as is or be used to explore climate change scenarios, and experiment with parameter variations. To do this, workflows are defined along multiple dimensions: (i) **Basin** [*Northern, Southern, Whole-*of-Basin], (ii) **Reference Scenario** [Without Development (*WoD)*, Basin Delivery Limits (*BDL*), Basin Plan 2024

(*BP2024*), Basin Plan Full Implementation (*BPFI*) (iii) **Connectivity** [*Disconnected*, up- to downstream *FlowLink*, *Feedback* (1, 3 or 5 iterations)] and (iv) **Functionality** [*Default, Climate Scaling, Experiment* (parameterization variations)], see Figure 6.

The Reference Scenario influences which models are part of the model chain, while the Connectivity dimension governs which data exchanges are applied in this model chain. The Functionality dimension determines whether additional preprocessing is required. In the GUI (Figure 7) separate workflow execution points are defined for the Basin, Reference Scenario and Functionality dimensions, while the Connectivity method and the valley model version to use are selected by the user in the FIRM Scenario. A selection in one of the dimensions will influence the available options in the remaining dimensions. For example, fewer models exist in a WoD Reference Scenario which has an impact on Connectivity options.

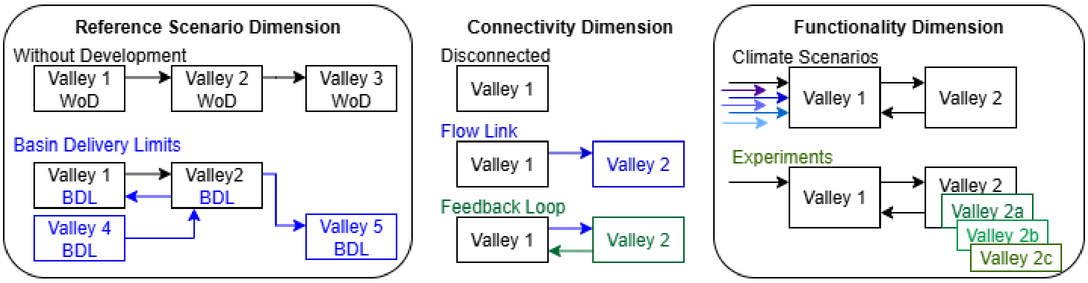

Figure 6 Three dimensions of workflows (left to right): Reference Scenarios, Connectivity and Functionality

At the time of writing, the system holds 88 model datasets executed by 26 different model code versions, defined for 4 different Reference Scenarios, 6 climate change scenarios and limitless options to experiment. Based on the selections made, the FIRM workflow ensures the appropriate models are executed and the relevant model linkages are applied. All workflows use the same timeseries administration (Valley, Location, Parameter, Iteration index, Reference Scenario and Functionality), while the run specific settings (which workflow, model versions, connectivity etc) are stored in the run's meta data (i.e. Scenario definition). This flexible structure allows for comparison of results between any and all scenarios and provides full traceability of the choices made in a scenario.

2.5. FIRM application

The compute and analysis environment is provided by the **FIRM application** (based on Delft-FEWS), the customised application using Delft-FEWS, a Java based generic data-model integration software platform mostly used in the real-time operational domain [Werner et al. 2012]. With the 'Computational Framework' developments by Deltares in 2020/2021, the scientific workflow engine limitations identified [Gijsbers, 2010] have been overcome, which makes the software platform suitable for data model integration for the scenario analysis domain.

Delft-FEWS acts as a scientific workflow engine, orchestrating the sequence of valley models executed, administering and storing model results and exchanging timeseries between different valley models. As part of the workflow, statistical postprocessing is conducted for climate scenarios runs and experiment workflows also include preprocessing steps. The instructions for the Delft-FEWS software are encapsulated in its configuration, which is customized using workflow templates and model specific metadata. As the model results are the basis for policy making, long-term storage of results and its production method will be established using the Delft-FEWS Open Archive [Grijze et al., 2014]. The model data management both in FIRM as in the archive provides full transparency and traceability of inputs, model versions, outputs, and user actions (who, what, when).

The FIRM application is deployed in Azure, executes climate scenario and experiment model runs in parallel across multiple computational nodes and can be scaled up or down depending on demand. Users access FIRM through a Windows desktop application that is distributed through MDBA's company portal to the virtual workstations.

The FIRM application is the entry point for the modeller to compose a scenario, execute a workflow, and analyse model results and data transfers between valleys, see Figure 7.

- On the left side, the workflows are organized by Basin (root level) and Reference Scenario and Functionality (level 2). Scenarios created by the user are added to this list (level 3). As scenarios are added, the tree can be restructured as needed.
- The centre panel shows the FIRM Scenario composition panel with the option to select the Connectivity method as well as the model version(s), and associated meta data, used for this Reference Scenario. In case multiple versions exist, further selection needs to be made by the user.

- The filters on the right-hand side help to navigate the model results of each Scenario, which can be a quite large dataset, depending on the scenario. For example, a Scenario with a Feedback loop for the Whole of Basin produces 130,000 timeseries. Meta data manages which series are saved, see section 2.2.

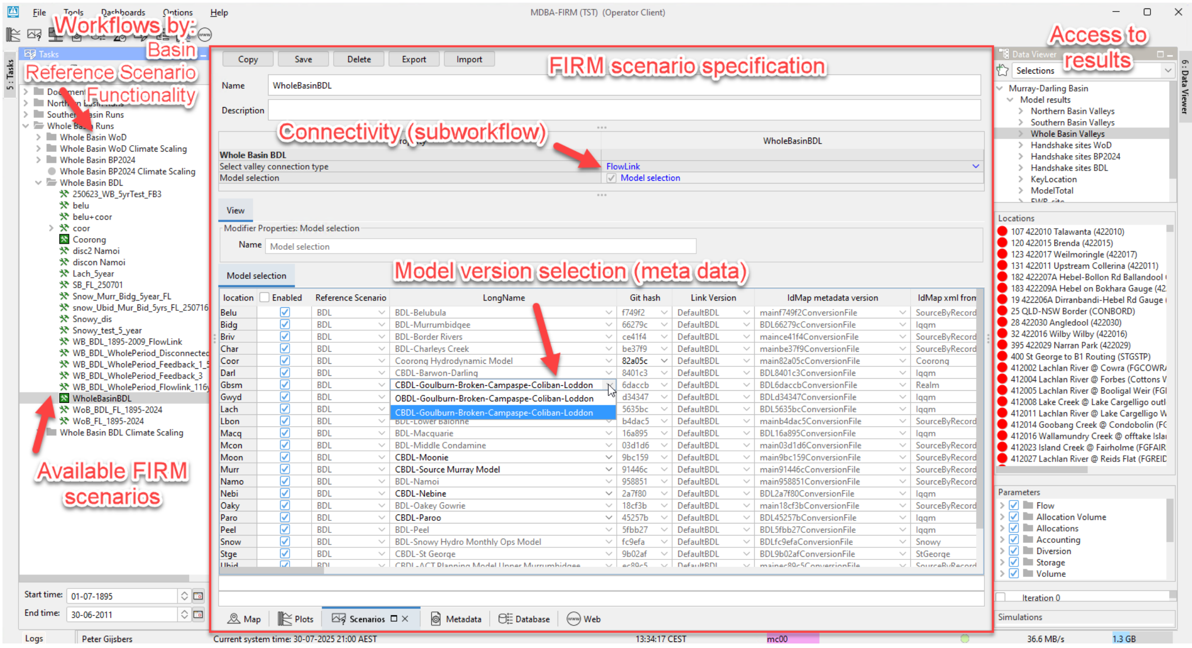

Figure 7 Scenario specification in the FIRM application

Logically, results visualization and comparison is an important part of FIRM. Detailed inspection and comparison of timeseries results is conducted using the FIRM application. For overview and a summary of the water balance, a separate dashboard web-application has been created with the name **FRAME**. FRAME obtains the necessary data to visualize the water balance in diagrams and tables directly from the FIRM application's database using the Delft-FEWS webservice API.

2.6. FRAME

While the FIRM application provides access to the detailed results of all model runs within the system, there is often a need for higher level visualisations that integrate multiple results relevant to a particular investigation, such as integrating all top-level water balance components for a valley. **FRAME** was built to provide these higher level visualisation products, both for regular users of FIRM and for other stakeholders who don't require access to the full set of results. FRAME is intended to grow along with the wider adoption of FIRM, with present functionality focussed on valley scale water balance reporting including waterfall charts of individual scenarios and various options for showing comparisons between scenarios (see Figure 8). A waterfall chart provides insight in the water balance by accumulating each flow component.

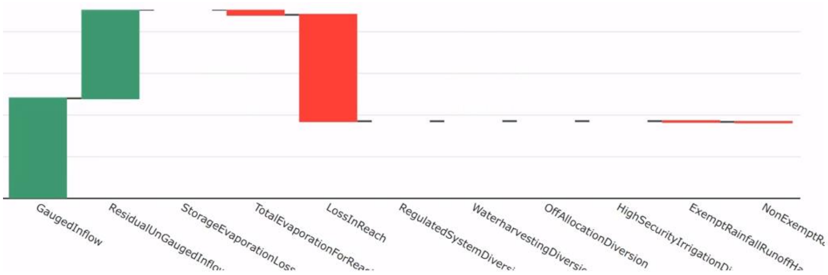

	FL 1895-2024	010880WoD
TotalStorageInReach	0.00 GL	0.00%
ExtraLossFromGwUse	6.45 GL	-8.16%
StorageVolume	14792.22 GL	-0.01%
TotalInflow	165.97 GL	[illegible]
GaugedInflow	58.67 GL	[illegible]
ResidualUnGaugedInflows	107.30 GL	[illegible]
Loss	13.59 GL	-17.34%
StorageEvaporationLoss	0.00 GL	0.00%
TotalEvaporationForReach	0.48 GL	13.91%
LossInReach	6.66 GL	-29.13%

Figure 8 Example views from FRAME including (a) waterfall chart and (b) scenario comparison table

FRAME is built in the Hydrograph visualisation platform, which provides a connection between a library of generic visualisation elements and a series of relevant data sources, including the Delft-FEWS API. FRAME accesses data directly from FIRM using this API, ensuring that FRAME is always able to display the latest data from the system. The entire Hydrograph dashboard is described in a JSON configuration file, including page layout, configuration of individual visualisation elements, user-interaction, data mapping and the calculation of derived data. The main Hydrograph application runs entirely within the browser and connects directly to relevant data sources, simplifying deployment by avoiding any additional component installation.

3. CHALLENGES FOR MODEL INTEGRATION

Model integration into a framework inherently poses specific requirements to the underlying modelling software. Compared to other Environmental Modelling Frameworks, FIRM poses few requirements as it runs models and exchanges data strictly in sequence, treating models as black box executables. As such, FIRM is not an invasive Environmental Modelling Framework [Lloyd et al. 2010], as it does not require API-callable libraries for a tight model integration.

However, even with loosely coupled model integration, challenges remain due to executable behaviour and coding methods applied. This may be expected for some of the legacy hydrological modelling software involved in FIRM which originate in the eighties, but even recent releases of more modern software do not always keep overarching system integration and headless server-side computation in mind. Some issues experienced during the implementation of FIRM were:

1. Command line calls which trigger Windows pop-ups that either require user interaction or pose the risk that some Windows-background service stall the process, bringing the entire workflow to a halt
2. Command line interface which lacks the *full* functionality to control the execution
3. Crashes when multiple models running the same software version in parallel on the same machine, e.g. because the software uses the same temporary file to manage locking
4. Frequent, non-backward compatible, changes to input and output file formats between software versions
5. Undocumented specific line ending needs (e.g. LF vs. CRLF)
6. Differences in modelling choices between Valleys, such as units (e.g. Salinity vs EC), daily vs monthly timesteps

4. CONCLUSIONS & OUTLOOK

In conclusion, FIRM represents a significant advancement in basin-scale hydrological modelling for the Murray–Darling Basin. By enabling the integration of diverse and legacy models within a cloud-native, metadata-driven environment, FIRM addresses long-standing challenges in collaboration, and transparency across jurisdictions, in a reproducible manner. The meta-data model at the heart of FIRM allows for flexibility in adding more models, for more reference scenarios, with new functionalities, making it a truly future proof collaboration platform. The combined power of the framework components Delft-FEWS, FOM, and FRAME, allows for flexible scenario analysis and supports policy-relevant modelling efforts.

While technical and interoperability challenges remain, FIRM lays the groundwork for a more unified and efficient approach to water resource modelling in Australia. Continued development and engagement with its users will be essential to fully realize its potential and ensure its long-term sustainability and impact.

REFERENCES

Brenan E, Peralta L, Vorreiter D, Fuller J, Price N, Dixon S, Gijsbers P, De Kleermaeker S, Cressall, B, Urich C, Rahman J (2025) Shifting Left: Building a collaborative integrated hydrological model for the Murray-Darling Basin The 26th International Congress on Modelling and Simulation (MODSIM 2025)

Gijsbers P (2010) Opportunities and limitations of Delft-FEWS as a scientific workflow tool for environmental modelling. International Congress on Environmental Modelling and Software. 17. https://scholarsarchive.byu.edu/iemssconference/2010/all/17

Grijze A, Gijsbers P, de Rooij E, Pelgrim E, Van Den Akker O (2014) Technology behind the Deltares Open Archive. International Congress on Environmental Modelling and Software. 16. https://scholarsarchive.byu.edu/iemssconference/2014/Stream-A/16

Lloyd W, David O, Ascough II JC, Rojas KW, Carlson JR, Leavesley GH, Krause P, Green TR, Ahuja LR (2010) Environmental Modeling Framework Invasiveness: Analysis and Implications. International Congress on Environmental Modelling and Software. 577. https://scholarsarchive.byu.edu/iemssconference/2010/all/577

Murray-Darling Basin Authority (2022) $66 million modelling upgrades to bring better Basin information, Murray-Darling Basin Authority, $66 million modelling upgrades to bring better Basin information | Murray–Darling Basin Authority, Accessed: 7 August 2025.

Werner M, Schellekens J, Gijsbers P, van Dijk M, van den Akker O, Heynert K (2013) The Delft-FEWS flow forecasting system. Environmental Modelling & Software 40(65-77) https://doi.org/10.1016/j.envsoft.2012.07.010.

Yang A (2010) The Integrated River System Modelling Framework. A report to the Murray-Darling Basin Authority. CSIRO: Water for a Healthy Country National Research Flagship, Canberra, June 2010

