

OPENGRID PLATFORM COMPONENTS – WORKFLOW ORCHESTRATION

One open platform, powerful interoperable components

OpenGrid's foundation is robust, trusted, open data, with task-ready applications layered on top. Each component was shaped by conversations with the grid planning community and user research on modeling pain points, and each deep-dive sets out what we heard, the outcome we're working toward, and a draft product concept for the community to react to.

Each deep-dive draws on more than 65 conversations with over 55 participants globally across the grid planning ecosystem, conducted in early 2026. We group participants into six personas along two lines: whether they produce analyses or use them, and whether they sit inside or outside the organization whose decisions are at stake. Each group brings different needs, and together they shape how every OpenGrid component is designed. To protect confidentiality, interviewees are identified only by organization type. The draft product concepts are starting points for OpenGrid's Technical Committee and working groups to refine, and we expect scope, priorities, and naming to change as that work happens.

This summary describes the current vision for each component of the OpenGrid Platform. They are not definitive specifications and are still subject to ongoing refinement. Each component will be shaped by the Technical Committee and a working group that is open to anyone who wants to take part. We expect scope, priorities and naming to change as that work happens. Comments, corrections and alternative proposals are welcome. We are publishing at this stage deliberately, so that direction can be shaped in the open rather than presented once it is already fixed.

Workflow Orchestration

Pain points, from user research

A grid study is a chain of data preparation, model runs, and reconciliation, and today that chain is stitched together by hand and held in specialists' heads. Nearly every distribution and transmission modeler raised it as a key barrier.

What we heard

- A US university said open-source tools exist for capacity expansion, production cost, transmission, and siting, but handoffs between them take weeks to months, and a single study can take a year.
- A US IOU reported iterations of 2 to 4 hours at best and up to 4 days at worst. About 70 percent of studies require extended time, and about 20 percent need a completely different approach mid-study.
- A European TSO has no formalized tool chain and relies on specialists who know which scripts to run. Its models carry 10,000 to 100,000 parameters, and quality checks are mostly manual. The hardest part is the "glue code": unstandardized scripts that turn raw data into model inputs.
- Another European TSO runs a four-step chain (scenarios, demand modeling, market simulation, grid calculations), and a vendor switch meant rebuilding every workflow from scratch.
- A European DSO handles about 100,000 connection requests a year in one region through a multi-step process, and material cost changes can invalidate completed analyses and force time-intensive reruns.
- A California CCA spends several weeks manually combining vendor forecasts with its stochastic uncertainty tool, and repeats the whole process every 4 to 6 months when new forecasts arrive.

- A US municipal utility said codebases become so bespoke that keeping them organized enough for others to use is a major barrier, and Python environment management is painful even for technically savvy users.

What users asked for

- An open resource planning initiative wants round-trip iteration between capacity expansion and production cost.
- A US ISO wants to streamline the chain from data collection to model building to study execution.

Vision

Any planning study can be rerun from raw data to final results, by someone other than the person who built it, and produce the same answer.

The weeks lost to moving data between steps, repairing handoffs, and remembering which scripts run in which order give way to workflows that are defined once, shared, and reused. Knowledge of how a study is done lives in the workflow itself, so it survives staff turnover and can be picked up by a colleague, a regulator, or a partner in another country. Proven, industry-grade workflows are available as starting points, which matters most where planning capacity is still being built.

Workflows can be assembled by people or by AI agents, and either way a person can see what ran, with which inputs, and why. Swapping one tool in a workflow means changing one step while everything around it stays in place.

Draft product concept (for discussion)

Workflow Orchestration is the layer that connects the others into a complete study. A modeler describes a multi-model analysis once, as a versioned workflow definition that names the kinds of models involved, how data moves between them, and the logic that governs the run, then reruns it with automatic, schema-aware handoffs between steps. Definitions can be written in a structured text format or built programmatically, and shared and registered so others can find and reuse them.

Real planning studies are rarely a simple chain, so workflows would support multiple analytical patterns that can be combined within one study:

- 1. Sequential.** Each model runs to completion and passes its outputs to the next, as when capacity expansion results feed production cost and then power flow.
- 2. Iterative.** Models exchange results in a loop until they converge, such as round trips between capacity expansion and production cost.
- 3. Co-simulation.** Models step through time together and exchange data at each step, as in transmission-distribution co-planning.
- 4. Co-optimization.** Several models are solved as one coordinated optimization problem.

Around those patterns, workflow orchestration could include:

- **Swappable model blocks.** Each step names the type of model it needs, such as capacity expansion, production cost, distribution, or resource adequacy, so any tool that meets that interface can be swapped in as a configuration change. These interfaces build on Open Standards & Schema and Model Plugins.
- **Handoff helpers** for the transformations that make handoffs hard, such as aggregating and disaggregating between zonal and nodal models, or getting AC power flow to converge on production cost results.
- **Provenance and validation.** Definitions and handoffs are checked before a long run starts, and each run records model versions, inputs, data transformations, logs, and intermediate results, so any result can be traced and reproduced.

- **Change propagation** reruns the affected steps automatically when an assumption changes, saving days of manual re-entry.
- **Templates** for common studies, such as integrated resource planning, resource adequacy, transmission-distribution co-planning, and interconnection studies.
- **People and agents.** An analyst can configure, launch, and monitor a workflow, or a script or AI agent can compose and invoke one through an API. Either way, a deterministic engine executes the plan, so the process stays reproducible and observable.
- **Run anywhere.** The same workflow runs locally, in the cloud, or across both, so sensitive steps can stay on premises.

Where mature engines already exist for a pattern, the layer should build on them behind a common interface.

Boundaries. Orchestration coordinates steps that other components perform. It does not convert formats, execute solves, store data, or render results.

Who We Heard From

Our platform components were shaped by more than 65 conversations with over 55 participants across the grid planning ecosystem. We group them into six personas along two lines: whether they produce analyses or use them, and whether they sit inside or outside the organization whose decisions are at stake.

Producer Side — People Who Run the Models

Transmission-System Users. Engineers and planners at ISOs, RTOs, TSOs and vertically integrated utilities who model the bulk system, often chaining three to five specialized tools through a single study.

Primary Distribution-System Users. Operators and modelers of distribution networks, from Europe's largest DSOs to small municipal utilities with one- to five-person teams.

Point-Source Modelers. Developers, hyperscaler procurement teams and community energy buyers who model a single point of grid interaction, with analyses that must be defensible enough to underwrite capital commitments.

Consumer Side — People Who Use the Results

Internal Decisionmakers. People who use model outputs to allocate their own organization's capital and must trust analysis they did not run.

External Evaluators. Regulators, intervenors, advocacy organizations, think tanks and data-infrastructure bodies who use outputs to scrutinize decisions made by better-resourced entities.

Straddling Both

External Primary Modelers. Analysts at public utility commissions, reliability councils, think tanks and capacity-building organizations who run their own models to examine decisions made by others.