

OPENGRID PLATFORM COMPONENTS – COMPUTE & SOLVER

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.

Compute & Solver

Pain points, from user research

While solver performance is a barrier, the cost of solving is also often human time: runs that fail without explanation, long runs that must restart, and settings only a few specialists know how to tune.

What we heard

- A US RTO builds about 35 base power flow models a year, and a team of 5 spends 400+ hours a year just solving them.
- A US ISO runs production cost models that take about 12 to 16 hours even in the cloud; if errors occur, the entire run is redone. Its reliability tools lack built-in parallel processing, so it relies on cloud runs of up to 30,000 nodes at additional cost.
- The same ISO said open-source tools can't yet model systems as large as the Eastern Interconnection, and computational efficiency is the main barrier to wider adoption.
- A European TSO described a solver reliability gap (licensed software always terminates, while open source needs more careful data prep) and a gap in staff knowledge of solver mechanics, even among teams with many PhDs.
- A US municipal utility runs Pyomo with the open-source HiGHS solver and reports HiGHS has improved significantly. It chose Pyomo for better infeasibility diagnostics and said the industry uses unnecessarily tight tolerances for planning questions.
- A US state PUC found the utility's production cost model takes about a week to run, impractical to reproduce, so staff simplified models to run on laptops: "very hard from a regulatory standpoint to say our results are more reliable than what the company is presenting."
- A European DSO's staff use its primary planning tool tens of thousands of times a day, and cost data changes can double that as analyses are rerun, straining the cloud-hosted database behind it.
- A South Asian think tank reported a ten-year run in an open-source capacity expansion tool taking roughly 30 hours.

What users asked for

- A European TSO and a US municipal utility want solvers that terminate predictably and explain infeasibility.
- A US state PUC needs enough compute to reproduce a utility's study and avoid having to simplify it.
- A US ISO wants open tools that scale to interconnection-sized systems.

Vision

Open solvers are easy to use and fast enough to build and test industry-grade planning models, and computing power no longer decides who gets to participate.

Open-source solvers handle problems at the scale of real interconnections and behave predictably: they finish, and when a problem can't be solved they explain why in terms a modeler can act on. They are fast enough for the iterative work of building, testing, and exploring a model, and in some settings for an entire study. Where a commercial solver is the better choice for final runs, moving to it takes no rework.

Modelers move freely between their own machines and the cloud, building and testing where resources are plentiful and running sensitive cases wherever their data rules require. Industry-grade compute is available to regulators, researchers, and civil society, so a regulator can rerun a utility's full study. New methods, including AI-assisted ones, keep pushing solve times down across all solvers.

Draft product concept (for discussion)

Compute & Solver strengthens the open solver stack and reduces the human time spent getting runs to work. It focuses on making open solvers the well-supported default for building and testing models, with a clean path to commercial solvers where they are the better choice.

It could include:

- **Solver choice without rework.** First-class support for leading open solvers across problem types, such as HiGHS, IPOPT, and SCIP, and a clean way to swap in a commercial solver for final runs.
- **Diagnostics.** Infeasibility explanations that name the conflicting constraints, convergence troubleshooting, and mid-run progress monitoring so long runs stop being opaque.
- **Sensible defaults.** Tolerances that are easy to loosen for fast approximate runs, automated variable scaling, and solver configurations that are reproducible and shareable.
- **An execution layer.** Solves run locally or in the cloud without bespoke setup, with support for running sensitive cases wherever data rules require.
- **Access.** Compute support for under-resourced users such as regulators and civil society organizations.

Boundaries. This layer executes models; it does not formulate them (Model Plugins, Workflow Orchestration), own the data, or render results (Output Visualization). The focus is on building on existing open solvers, including new methods such as AI-assisted techniques that speed them up.

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.