Sandia National Laboratories

Exceptional service in the national interest

SAND2026-24275C

QUEST 3.0– PATHWAY TO AUTOMATIC WORKFLOWS FOR ENERGY STORAGE ANALYTICS

Presented by Tu Nguyen

U.S. DEPARTMENT of ENERGY | NNSA National Nuclear Security Administration

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

PROJECT OVERVIEW

- **Goal:** Advance QuESt from a collection of energy-storage tools into a workflow-centered analytics platform. The QuESt Workspace enables users to build executable visual workflows, while the QuESt AI Agent helps users understand, edit, extend, troubleshoot, and reuse those workflows.
- **Current Practice:** Energy-storage analyses are often split across separate apps, scripts, notebooks, manual data-prep steps, and result files. This makes studies harder to reproduce, share, modify, and connect across domains such as BTM, valuation, planning, performance, microgrid, and data management.
- **Why SNL:** Sandia is well positioned because QuESt already combines open-source energy-storage analytics tools, domain models, Python-based extensibility, and a user community focused on storage planning, valuation, performance, and grid applications. The Workspace and Agent build directly on that existing QuESt ecosystem.
- **Innovation:** QuESt 3.0 adds three major updates:
  - **QuESt Workspace:** Expanded into a richer canvas-centered workflow environment for master flows and subflows. Users can assemble data nodes, Python nodes, notebooks, subflows, and app-specific processing steps; manage Python environments and flow inputs; save/reload workflow files; run configured cases; and inspect generated outputs.
  - **QuESt AI Agent:** Now works directly with Workspace context. It can review the active canvas, use pinned messages and attached files, identify relevant tools and reusable skills, draft workflow action plans, preview changes before applying them, troubleshoot node outputs and run results, and support different models from OpenAI, Anthropic, Google.
  - **Installation Agent:** simplified deployment and configuration of open-source software, reducing technical barriers for users and accelerating adoption.

PROJECT OVERVIEW

- **Duration: FY26**.
- **Impact:** Researchers, analysts, utilities, DOE programs, and storage technology stakeholders benefit from more transparent, repeatable, and extensible energy-storage studies. Successful deployment reduces workflow friction, preserves expert analysis patterns as reusable skills, lowers the barrier for non-expert users, and makes cross-tool studies easier to inspect, rerun, and share.
- **Vertical/Investment Area:** Analytics
- **Alignment/Supports strategic pathway:** Supports DOE OE goals by improving tools for storage analytics, grid planning, and operational decision support. The Workspace helps stabilize and optimize grid services by making storage studies reproducible and connected; the AI Agent helps grow analytical capacity by guiding users through complex QuESt workflows and preserving reusable expertise.

Overview of QuESt Platform

QuESt has grown into a broader Python-based platform that hosts multiple QuESt tools and libraries for energy storage and related analysis. In addition to individual tools, QuESt includes a Workspace environment where users can assemble workflows that connect multiple tools, intermediate data preparation steps, and outputs into a larger study process.

QuESt App Hub

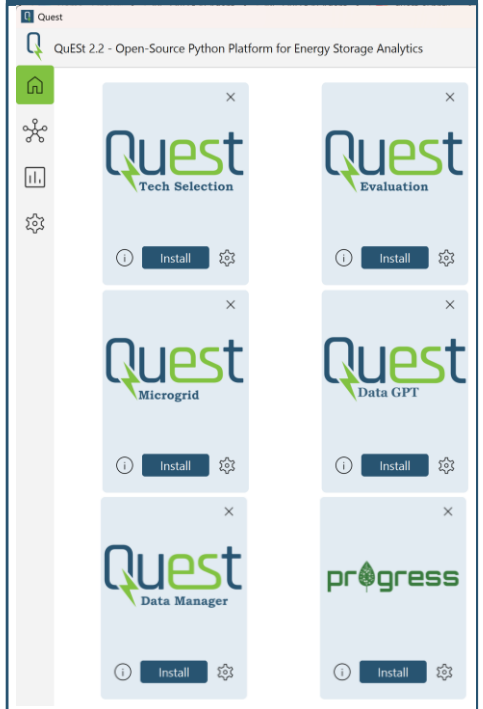

QuESt Workspace (Workflow Environment)

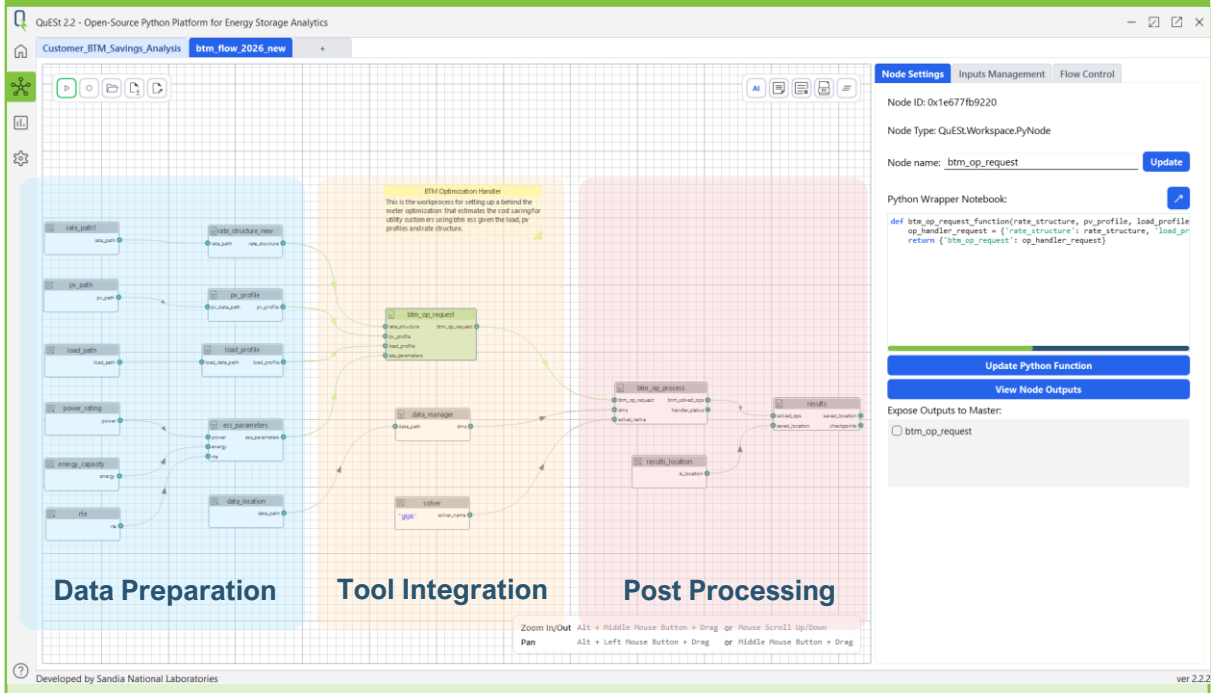

- **Connect Tools & Data**
- **Repeatable Processes**
- **Coordinate Multi Users**
- **Generate Outputs**

Challenges

Usability

- Which tool is appropriate?
- How to prepare inputs?
- How to configure nodes and flows?
- How to operate in a Python environment?

For Technical Users

Tedious but manageable with Python knowledge and experience

For Simple Users

Can be a barrier with technical complexity and python requirement

Impact

QuESt's impact is still partly limited by the amount of technical and python knowledge required to use the platform effectively.

Why it matters?

Many real analyses are not single-tool tasks. They require a sequence of steps, iteration, and coordination across different capabilities. The Workspace gives QuESt a strong foundation for more advanced studies.

QuESt Agent Architecture

The architectural logic is to keep the agent useful without making it free-floating. A language model is good at interpreting messy user intent, but QuESt needs operations that are grounded in the actual Workspace state. The MCP layer provides that grounding by acting as a controlled service interface between agent reasoning and QuESt internals.

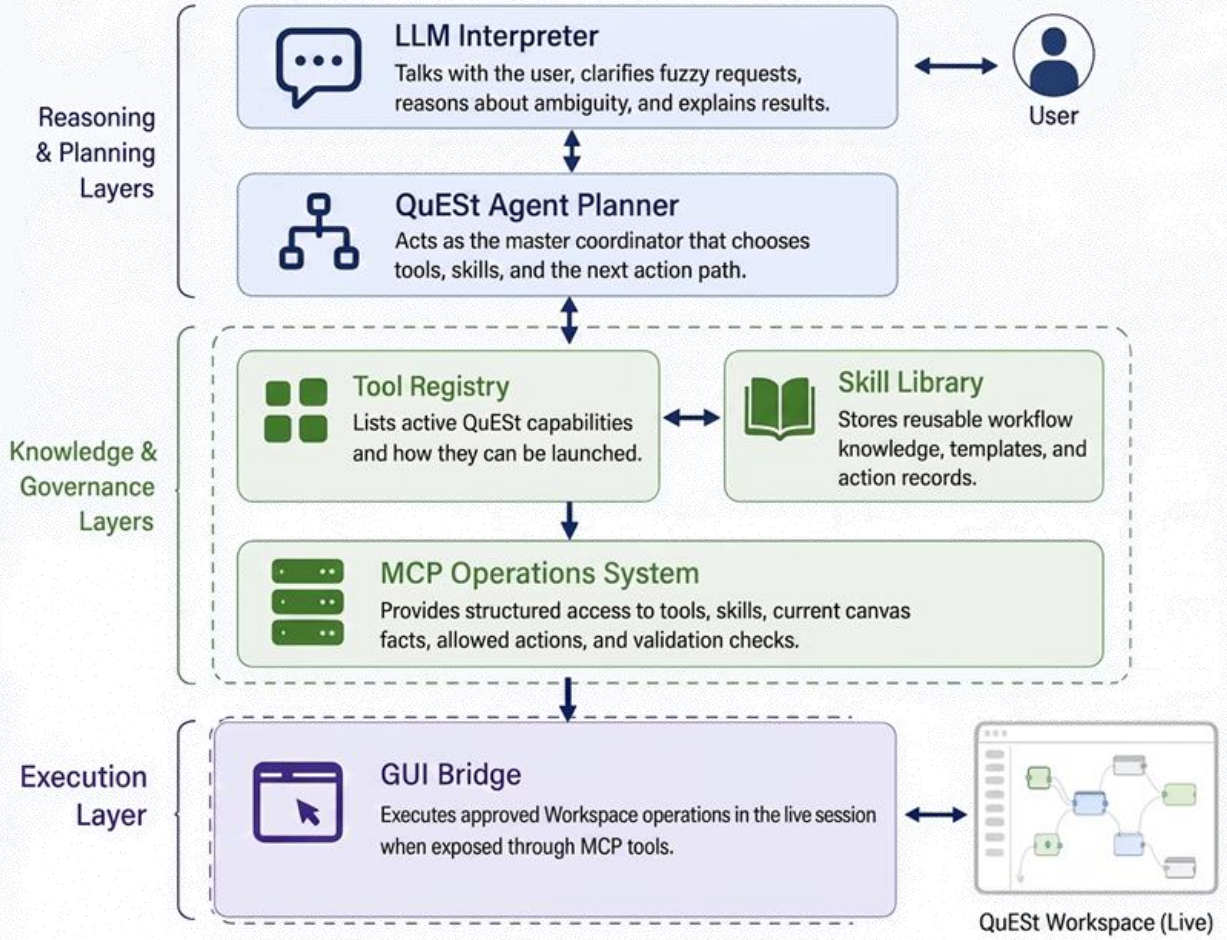

Layer	Role	Why it matters
LLM Interpreter	Talks with the user, clarifies fuzzy requests, reasons about ambiguity, and explains results.	Keeps natural-language interaction flexible without making the LLM the source of truth.
QuESt Agent Planner	Acts as the master coordinator that chooses tools, skills, and the next action path.	Keeps planning authority with QuESt Agent rather than with MCP alone.
Tool Registry	Lists active QuESt capabilities and how they can be launched.	Prevents tool selection from being guessed or hard-coded.
Skill Library	Stores reusable workflow knowledge, templates, and action records.	Turns prior workflow work into training material for future tasks.
MCP Operations System	Provides structured access to tools, skills, current canvas facts, allowed actions, and validation checks.	Gives the agent reliable inventory, routing, and quality-control facts.
GUI Bridge	Executes approved Workspace operations in the live session when exposed through MCP tools.	Separates planning from controlled application-side execution.

This work is supported by *Sandia's D8K AI Bigbet Program*

Mental Model: The QuESt Restaurant

A simple way to explain the MCP model is a restaurant. QuESt Workspace is the kitchen where flows are assembled, edited, inspected, and served. QuESt tools are the kitchen tools and stations, such as Workspace operations, Progress, Planning, BTM, and Python nodes. QuESt skills are recipes or cookbooks: general Python skills are basic recipes, while tool-specific skills are specialized recipes that require certain tools. QuESt flows are the dishes produced by the kitchen.

USER

Makes a request (messy or fuzzy)

FRONT OF HOUSE

LLM

Interpreter & Reasoning Assistant

- Talks with the user
- Interprets fuzzy requests
- Helps reason about ambiguity
- Explaine results

Role: Front-of-house interpreter and reasoning assistant

Conversation & Clarification

MASTER CHEF

QuESt Agent

Planner & Coordinator

- Plans the dish
- Chooses recipes and tools
- Coordinates actions
- Checks quality
- Decides what to do next

Role: Master chef that plans, coordinates, quality-checks, and drives execution

Access via MCP Operations System

RESTAURANT OPERATIONS SYSTEM

MCP

Operations System

Provides reliable access to:

- Tools (stations & equipment)
- Recipes (skills & cookbooks)
- Inventory (available resources)
- Kitchen state (current status)
- Routing (how to execute)
- Allowed actions
- Validation checks (quality control)

Role: Operations system that provides grounded access, routing, and validation

Directs the kitchen (through MCP)

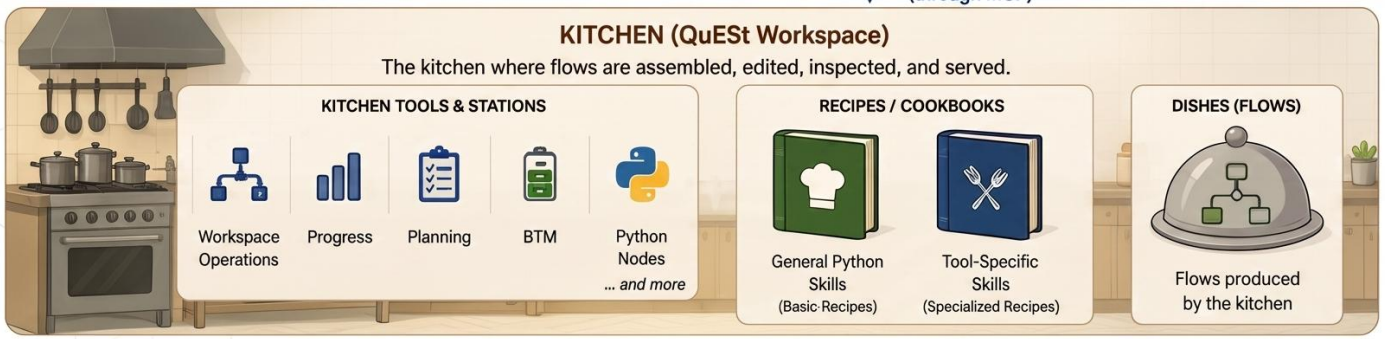

KITCHEN (QuESt Workspace)

The kitchen where flows are assembled, edited, inspected, and served.

KITCHEN TOOLS & STATIONS

- Workspace Operations
- Progress
- Planning
- BTM
- Python Nodes

... and more

RECIPES / COOKBOOKS

- General Python Skills (Basic Recipes)
- Tool-Specific Skills (Specialized Recipes)

DISHES (FLOWS)

Flows produced by the kitchen

KEY RELATIONSHIPS

QuESt Agent (Master Chef)
Has planning authority and decides what to do.

LLM (Front of House)
Interprets, reasons with the user, and explains.

MCP (Operations System)
Enables the chef with reliable, controlled access.

MCP is not the boss of QuESt Agent. It is the restaurant operations system that gives the chef reliable access to tools, recipes, inventory, kitchen state, routing, allowed actions, and validation checks.

Near-term Direction

1 The agent can be trained through use

Build or refine a flow, record the Workspace actions, preserve the workflow JSON and supporting artifacts, and promote the result into a reusable skill.

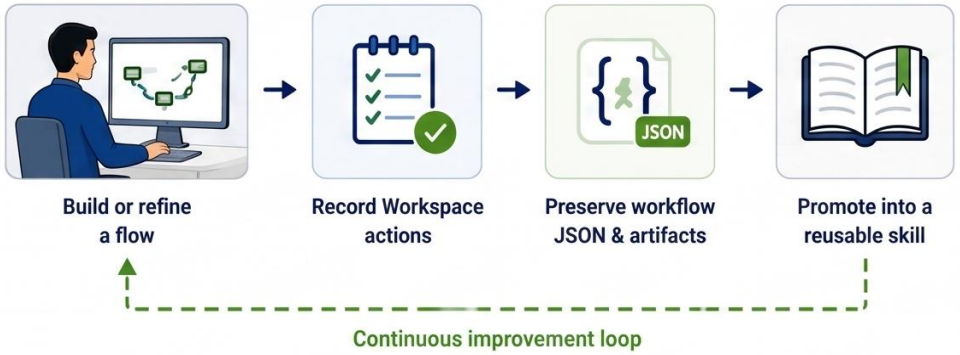

2 Future deep-agent orchestration can operate over explicit stages

The agent follows a clear sequence to reliably turn a request into an executed result.

1 Understand the request
Interpret user intent and clarify any ambiguity.

2 Match tools and skills
Find the right tools and skills that fit the task.

3 Inspect the current canvas
Examine the Workspace state, nodes, and data.

4 Validate missing pieces
Identify gaps and requirements to proceed.

5 Draft an operation plan
Create the step-by-step plan with tools, skills, and dependencies.

6 Execute through the bridge
Run the plan ria the MCP bridge and monitor results.

3 The same boundary can support both internal QuESt use and external MCP clients, making the architecture more maintainable as more tools and workflows are added.

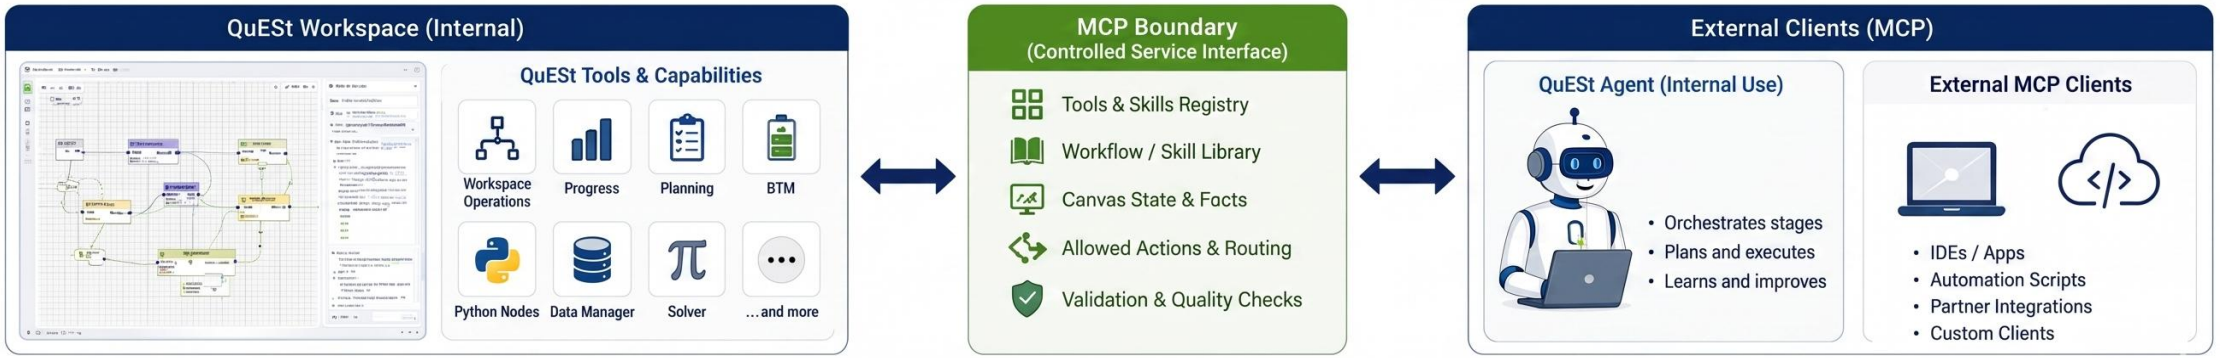

Outcome: A learn-by-doing agent, clear orchestration stages, and a stable MCP boundary enable reliable automation today and scalable intelligence tomorrow.

Long-term Direction

QuESt's vision is to be the one-stop shop for Sandia's Python tools and a catalyst for collaboration, technology transfer, and workforce development.

One-Stop Shop for Sandia's Python Tools

A unified platform where tools, data, workflows, and AI work together in one ecosystem.

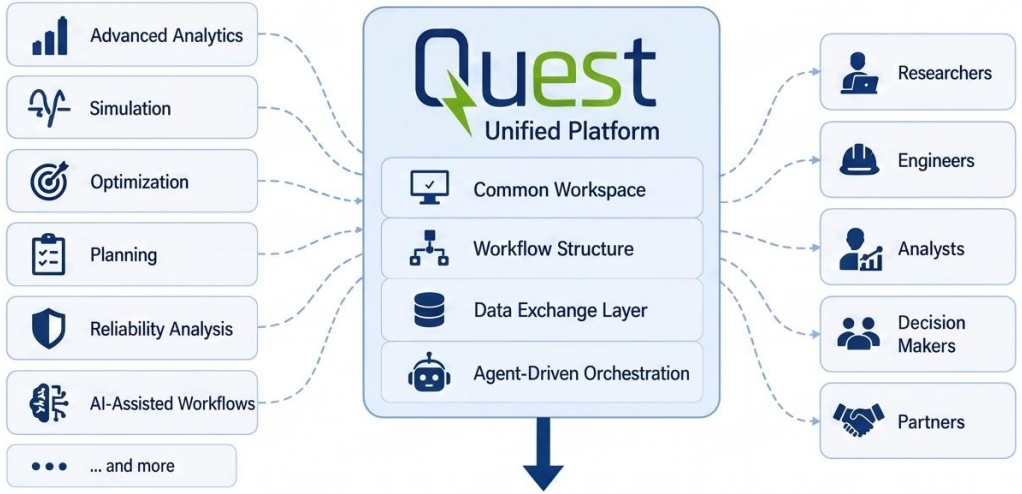

Lower Barriers
Easier onboarding, less technical friction

Higher Adoption
More users, more use cases

Greater Impact
Stronger tools, larger reach

Engagement & Ecosystem Growth

QuESt as a collaboration platform advancing Sandia's energy storage and grid modernization mission.

Proposals
Include QuESt in future proposal efforts

Workshops
Use in training, tech demos, and workshops

Universities
Collaborate on research and workforce development

Utilities
Partner with utilities on real-world challenges

National Labs & Industry
Strengthen ties and accelerate technology transfer

A Collaborative Ecosystem

Reusable workflows • Agent assistance • Standardized interfaces

Expanding access to advanced energy analytics and strengthening Sandia's impact

Outcome: A scalable platform and ecosystem that maximizes the value of Sandia's tools, drives engagement, and advances the future of energy and grid analytics.

DEMONSTRATION 1 – GRID RELIABILITY AND PLANNING INTEGRATION WITH QUEST AI AGENT

Team: Dilip Pandit, Atri Bera, Cody Newlun

- **Electric grid planning** is critical to tackle emerging challenges including **large load growth** and ensure a **secure**, **reliable** and **affordable** grid
- **Long-term electric grid planning** includes:
  - **Large** datasets and uncertain forecasts
  - **Intricate** coordination between complex models
  - **Sophisticated** solutions with long-term impacts
- Traditional processes are **manual** and **heavily time intensive**
- **QuESt Workspace** and **QuESt Grid Planning Toolbox** offer Sandia existing capabilities to strengthen and identify the role of AI

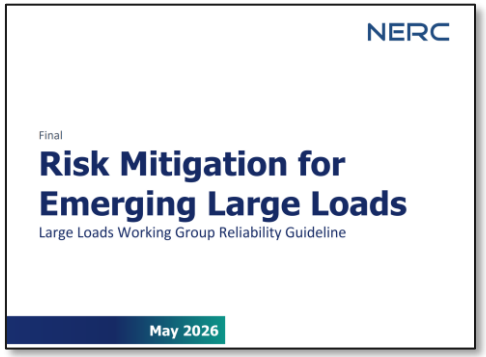

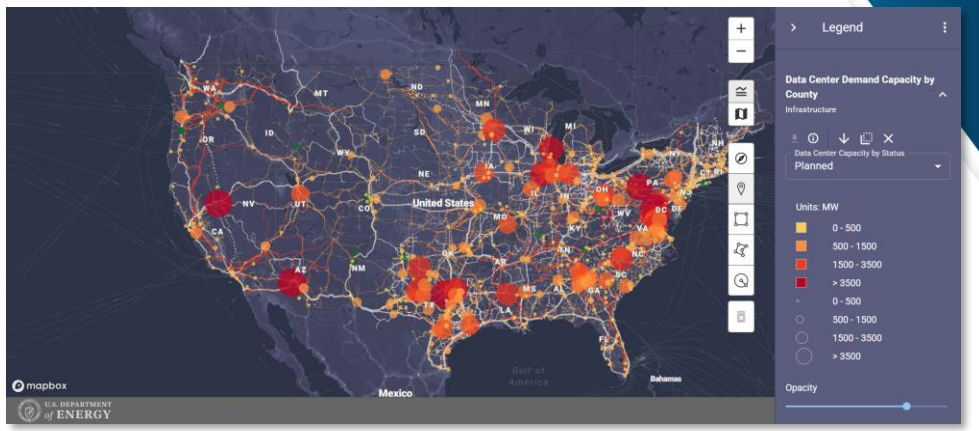

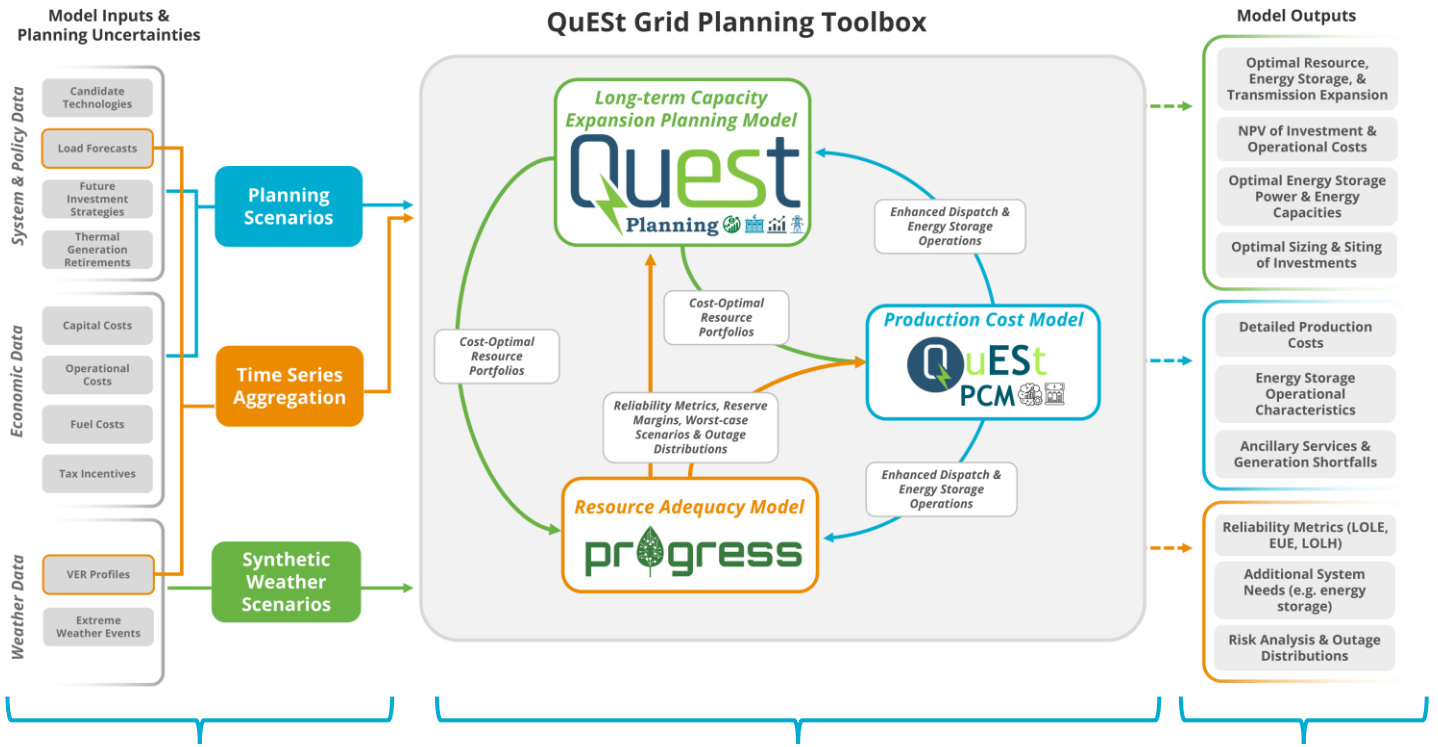

QUEST WORKSPACE FOR GRID PLANNING

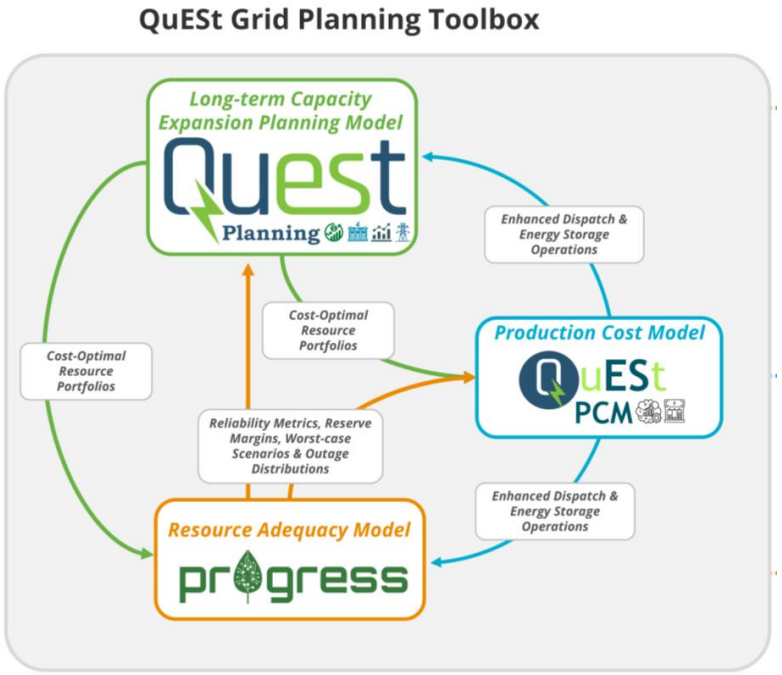

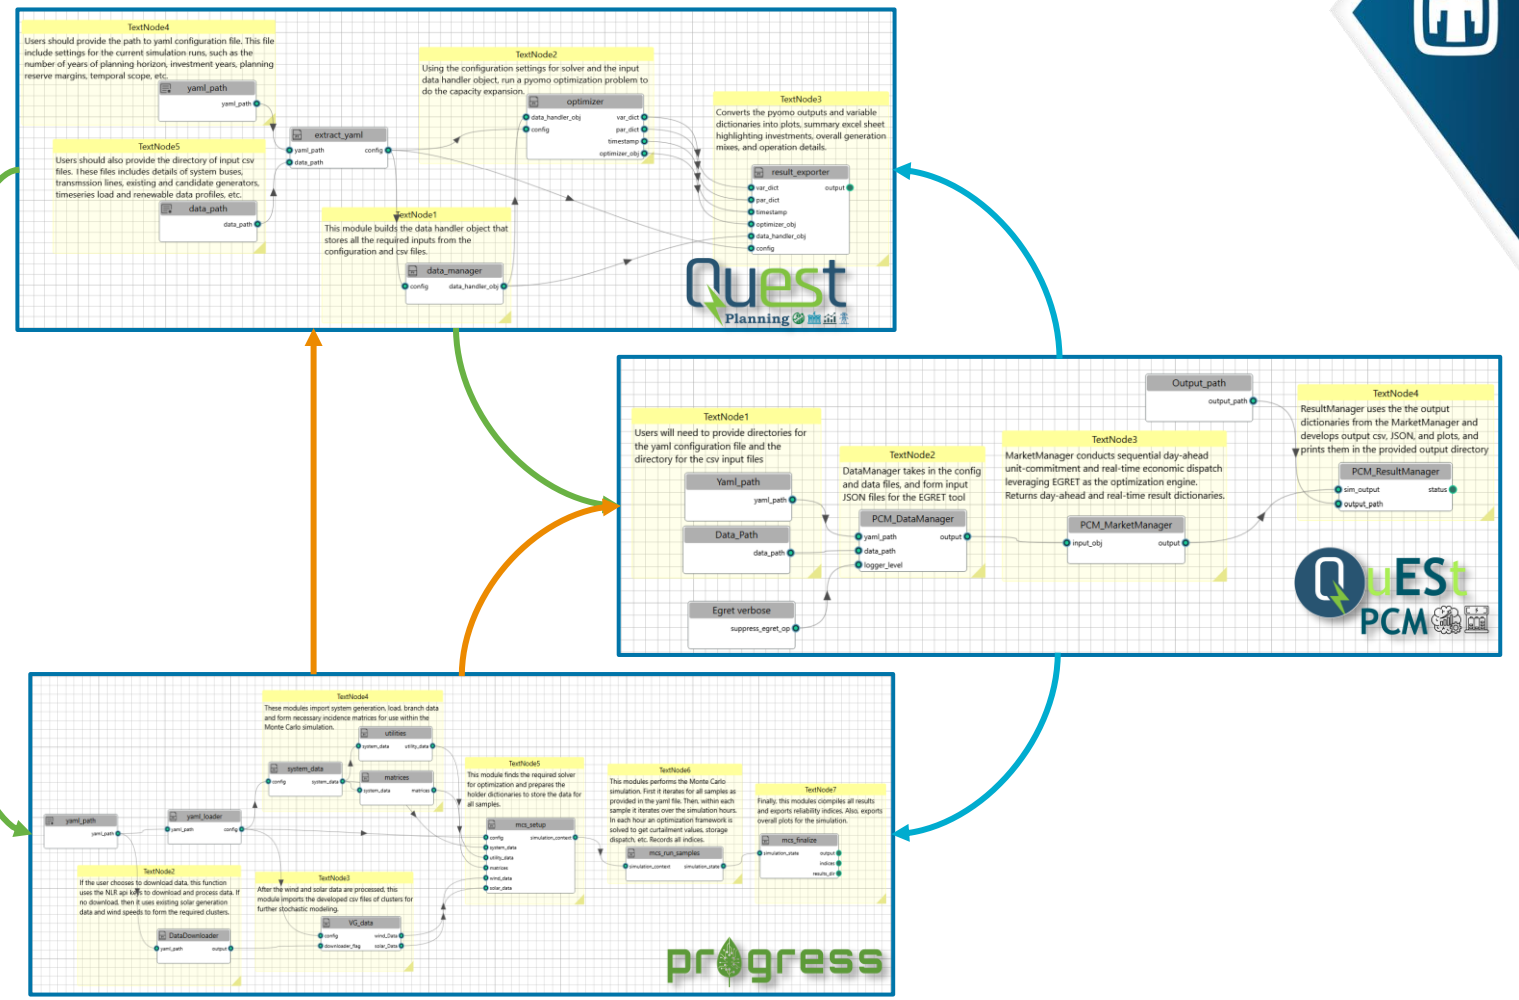

Tool Integration Challenges

- Tools run in different environments
- Simulation fidelity varies: decades for planning, yearly for adequacy, and sub-hourly for PCM
- Different data requirements for each tool

QuESt Workspace & Agents Innovation

- Integration of different Python environments
- Visual representation of modules, processes
- Enhanced workflow development, runtime debugging, and result interpretation

Training QuESt Agents

- Develop individual tool workflows using QuESt workspace
- Develop pipelines connecting the tools
- Develop skills for the agents on individual and combined workflows

QUEST AGENT FOR CO-PLANNING EXAMPLE

create workflow to do a power system capacity expansion and use it's outputs to perform reliability assessment on all investment years

GPT-5.4

1. User prompt

create workflow to do a power system capacity expansion and use it's outputs to perform reliability assessment on all investment years

10:12 AM

Build path: Reuse Skill Workflow JSON
Reason: Matched skill template 'IRP master flow for planning plus resource adequacy assessment' is an exact-enough fit, the canvas is empty, and no explicit differences were requested, so direct workflow reuse is the easiest path.

Proposed canvas plan:
1. Load the matched workflow template from skill `irp-master-flow-for-planning-plus-resource-adequacy-assessment`.
2. Validate current flow and prepare the next fix-up plan if anything is still missing.

Next step (1/2): Load the matched workflow template from skill `irp-master-flow-for-planning-plus-resource-adequacy-assessment`.
Use the actions below to execute, review, or cancel this plan. You can still send a revised request to change it.

Next | Show Plan | Cancel

2. QuESt Agent responds and matches closest skill

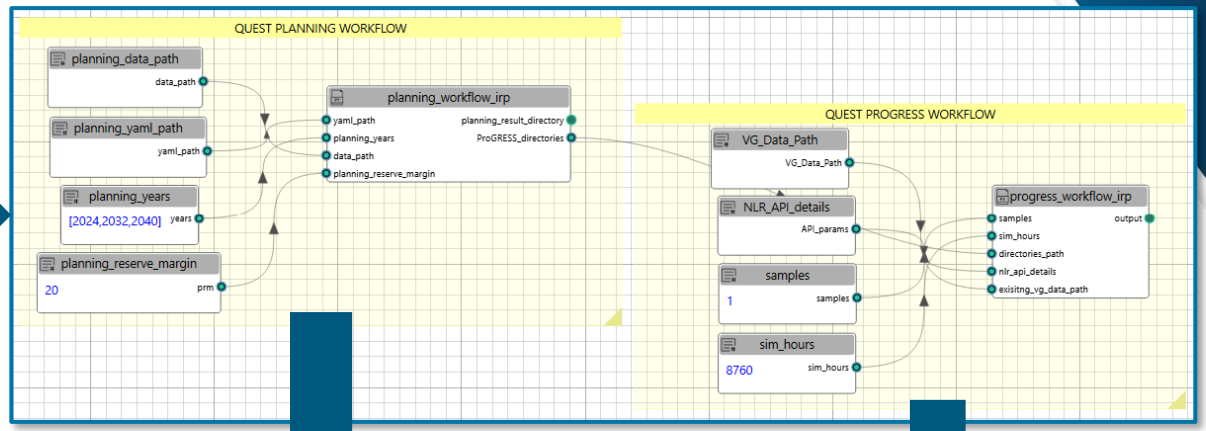

3. QuESt Agent loads workflow

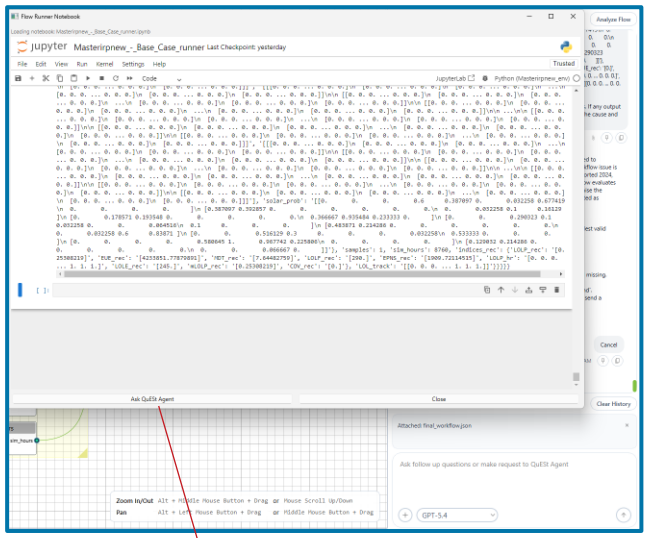

4. User asks Agent to interpret results

What the results say now
- 2024: Planning keeps a broad conventional fleet online, and Progress shows no adequacy issues.
- 2032: Planning has already retired a lot of coal and oil and leans more on renewables plus some storage; Progress still shows no shortfall in this test.
- 2040: Planning builds a very large renewable system, especially new solar and wind, with some storage, while conventional thermal is almost gone. In the exported 2040 Progress system, only 6 thermal generators remain plus storage. Progress then shows clear shortfall: LOLP 0.2, LOLH 20, EUE about 27,172, LOLF 3, LOLE 3.

How to interpret the investment-reliability link
- Planning is finding a least-cost expansion under its own assumptions: seasonal blocks, reserve constraints, transportation-style transmission, deterministic representation, and no annual load shedding cost.
- That solution is heavily renewable by 2040, with very little firm thermal left.
- Progress is then stress-testing that exported system stochastically hour by hour with outages and renewable variability.
- So the 2040 result means the least-cost Planning portfolio is likely not firm enough under the current Progress setup, especially once outages and chronological variability matter.

Why the 2040 result is suspicious but still informative
It is informative because it points to a real risk: the 2040 plan may be over-dependent on variable generation and too light on firm capacity.

But it is also suspicious because the two tools are not evaluating the exact same system model:
- Planning uses tx_model = Transportation with reserves enabled.
- Progress is forced to model = Copper Sheet.
- Planning solves on seasonal representative structure; Progress simulates chronological stochastic conditions.
- The exported Progress system is simplified CSVs, not the full Planning formulation.

5. Agent provides insights

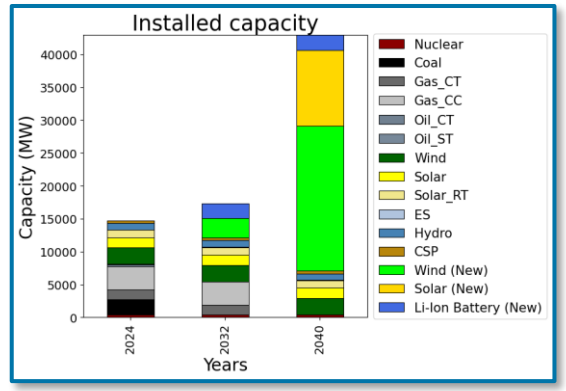

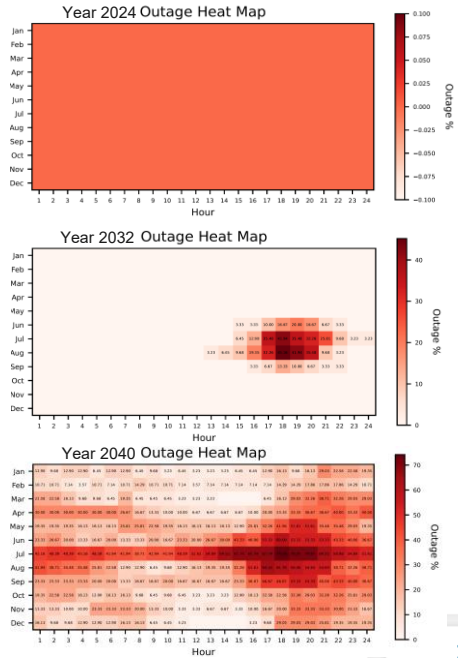

HOW AI MULTIPLIES THE IMPACT & NEXT STEPS

AI Agent Capabilities for Grid Planning Workflows

Impact: Enables users from non-domain backgrounds to directly import and apply standardized, rapid grid planning workflows

Next Step: Expand the skills for multiple QuESt tools and pipelines

Workflow Development & Customization

Impact: Users can develop and customize workflows for planning applications leveraging QuESt and other open-source tools

Next Step: Enable import of relevant python nodes from different sub-flows across different skill templates

AI-Assisted Engineering Analysis

Impact: Supports automated analysis of workflow outputs such as outage impact, adequacy metrics, and system reliability indices

Next Step: Automate the platform to implement suggested remedial actions (example: enable automatic tuning of planning model parameters)

QuESt Grid Planning Toolbox

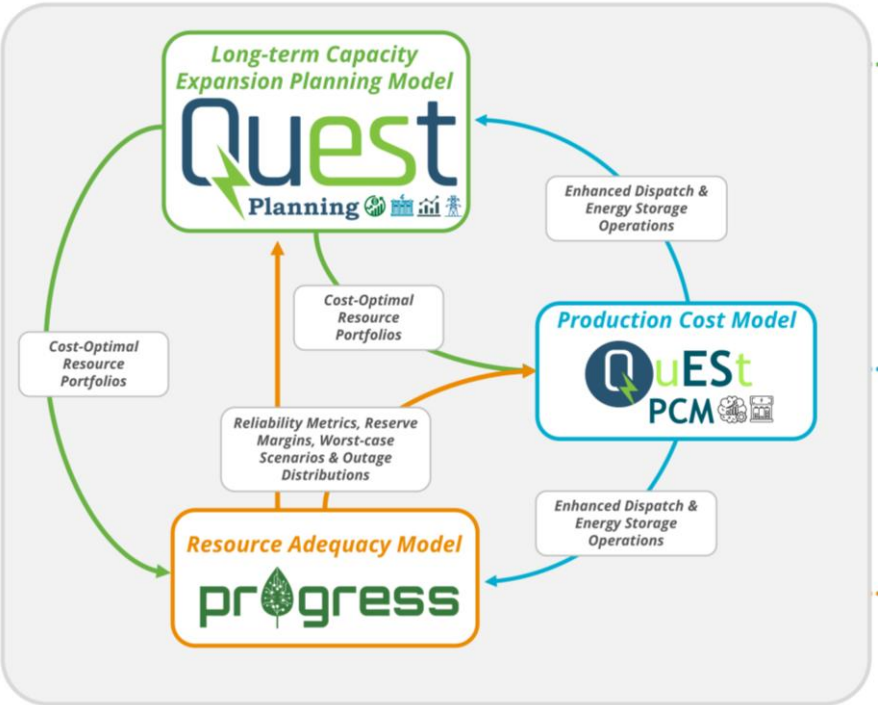

DEMONSTRATION 2 - BATTERY SIMULATION WORKFLOW WITH QUEST AI AGENT

Objective: Develop an AI-assisted battery simulation workflow within QuESt to analyze detailed electrochemical battery behavior across a range of QuESt scenarios and solution pathways.

Team: Hyungjin Choi, Tu Nguyen

- **Why is the work needed?**
  - QuESt battery techno-economic tools often use simplified battery models for computational efficiency, but these can miss key electrochemical dynamics such as degradation-driven capacity loss.
  - Physics-based model simulations often require specialized expertise, creating a high barrier to entry for general users.
- **What's new in our approach?**
  - We implement a QuESt workflow that uses PyBaMM to simulate heuristic/physics-based battery models for techno-economic scenarios and solutions.
  - We leverage the QuESt AI tool to streamline the configuration of required simulation setups, including the selection of appropriate physics-based models and parameters.

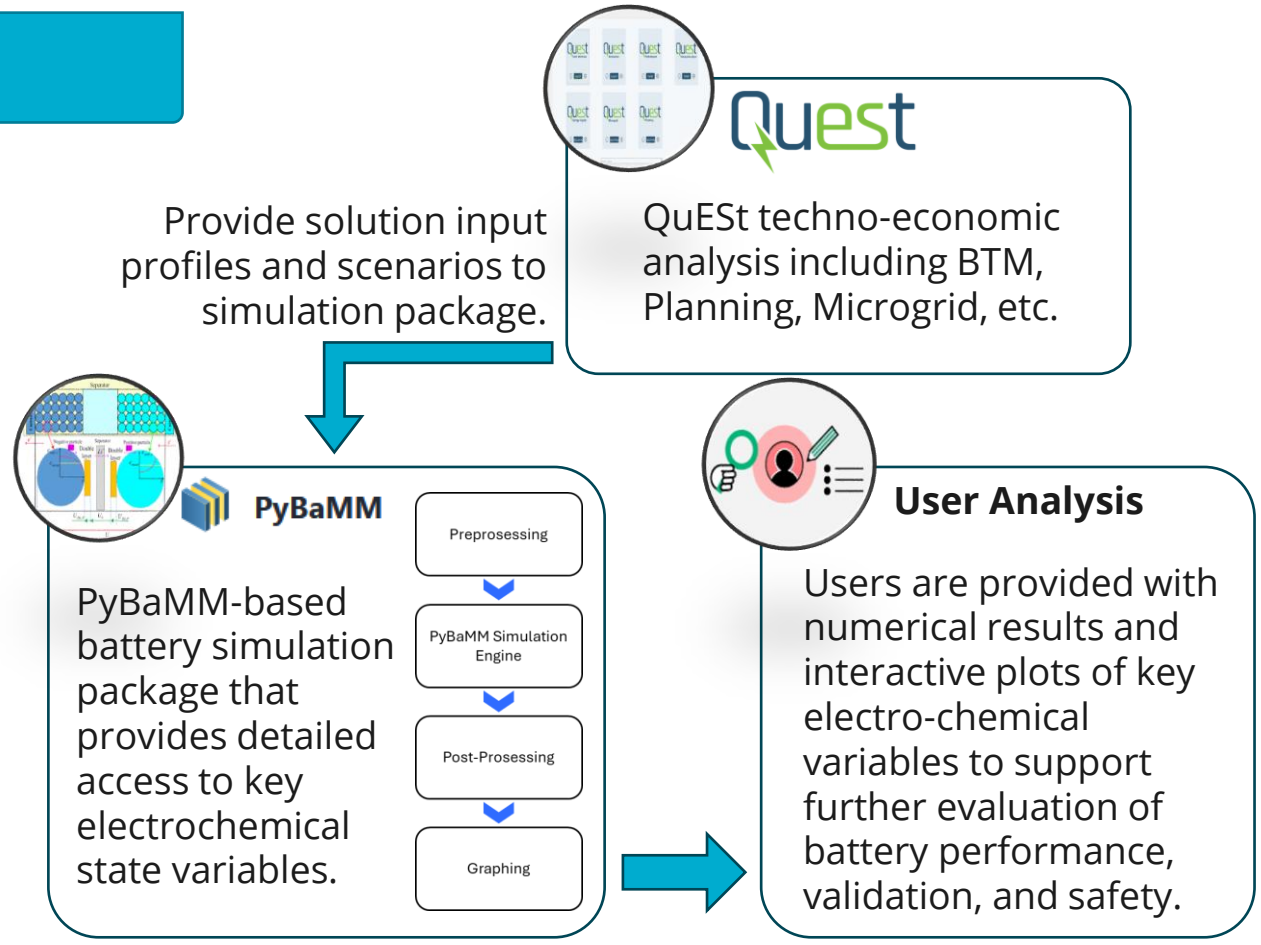

QUEST-PYBAMM SIMULATION WORKFLOW

Scenario loader
Load scenario files containing models, parameters, input profiles, economic variables.

Preprocessor
Convert scenario files into PyBaMM-compatible simulation inputs.

Simulation Engine
Simulate high-fidelity battery models.

Postprocessor
Process raw simulation data into user-friendly data sheets.

Graphing
Visualize results to compare scenarios and support analysis and validation.

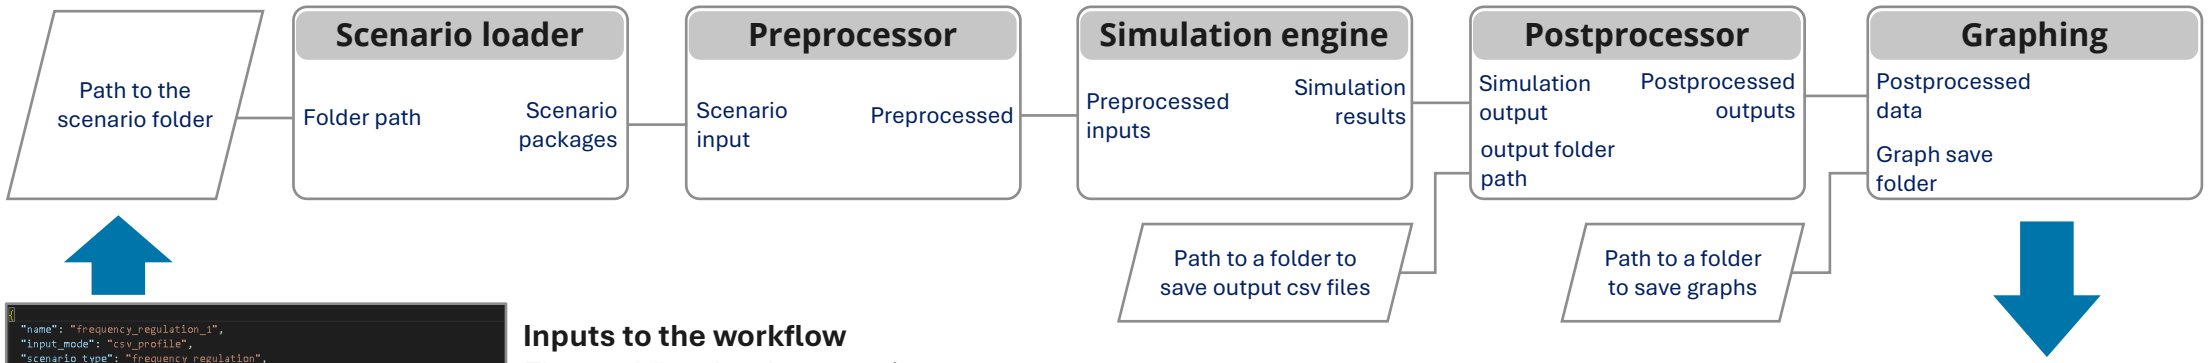

```
{
  "name": "frequency_regulation_1",
  "input_mode": "csv_profile",
  "scenario_type": "frequency_regulation",
  "model_type": "SPM",
  "parameter_set": "OKane2022",
  "battery_family": "Lithium-ion",
  "chemistry": "NMC",
  "cathode": "Nickel-manganese-cobalt cathode",
  "anode": "Graphite",
  "electrolyte": "Liquid electrolyte",
  "input_profile": {
    "csv_file": "W:/QuESt_AI/workflow profile data/frequency_regulation_1.csv",
    "type": "power",
    "units": "MW",
    "period": "60 seconds"
  },
  "model_options": {},
  "pack": {
    "cells_in_series": 15,
    "cells_in_parallel": 5,
    "profile_level": "pack"
  },
  "initial_state": {
    "initial_soc": 0.5
  },
  "capacity_loss_cost": {
    "cost_per_kWh": 150.0,
    "currency": "USD"
  },
  "solver": {
    "name": "CasadiSolver",
    "mode": "safe",
    "atol": 1e-06,
    "rtol": 1e-06
  }
}
```

Inputs to the workflow
The workflow loads scenario files that contain battery specifications and techno-economic variables to evaluate.

Outputs
The outputs include all numerical values as CSV files and interactive plots for all scenarios, enabling users to perform battery performance analysis and techno-economic evaluation.

Workflow implemented in QuESt workspace

WORKFLOW OUTPUTS

- **The simulation workflow generates detailed battery performance data, including current, voltage, power, state of charge, capacity loss, and energy throughput.**
- **Outputs are provided as CSV files containing numerical results, along with interactive plots for available internal battery states and key summary/statistical variables.**

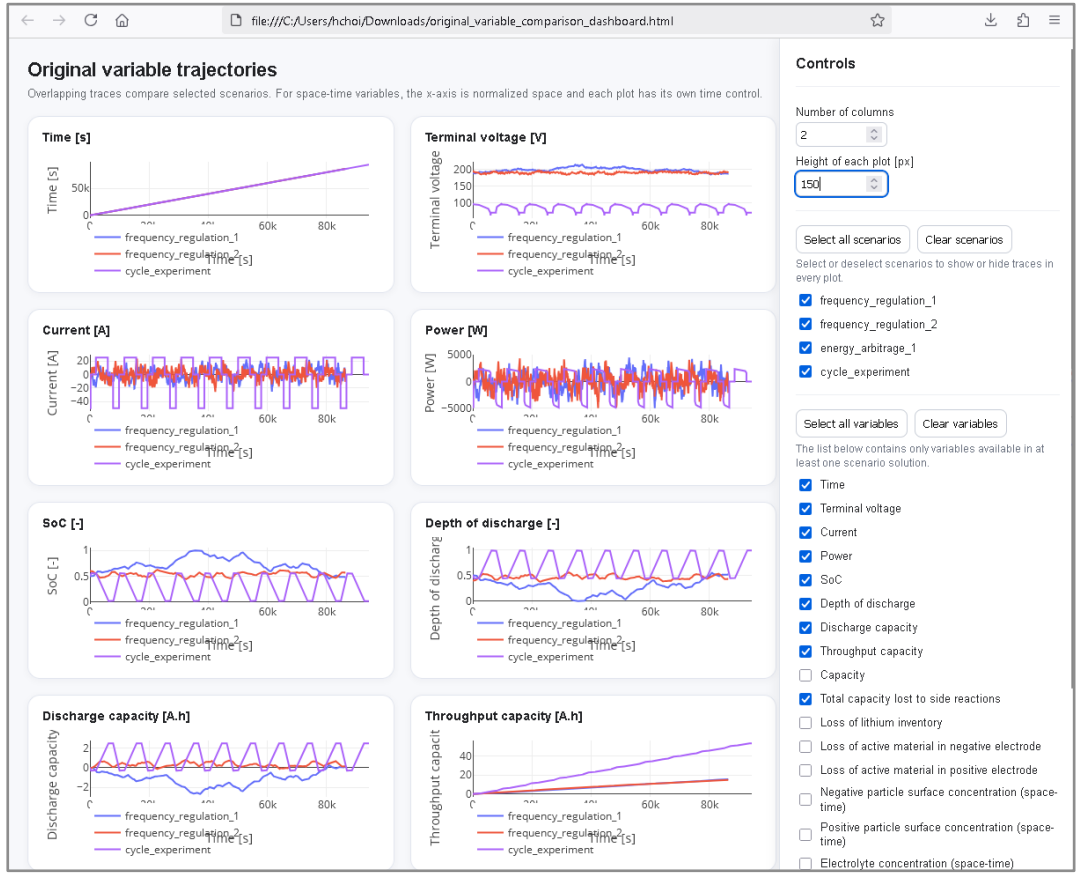

Battery electrochemical state plots

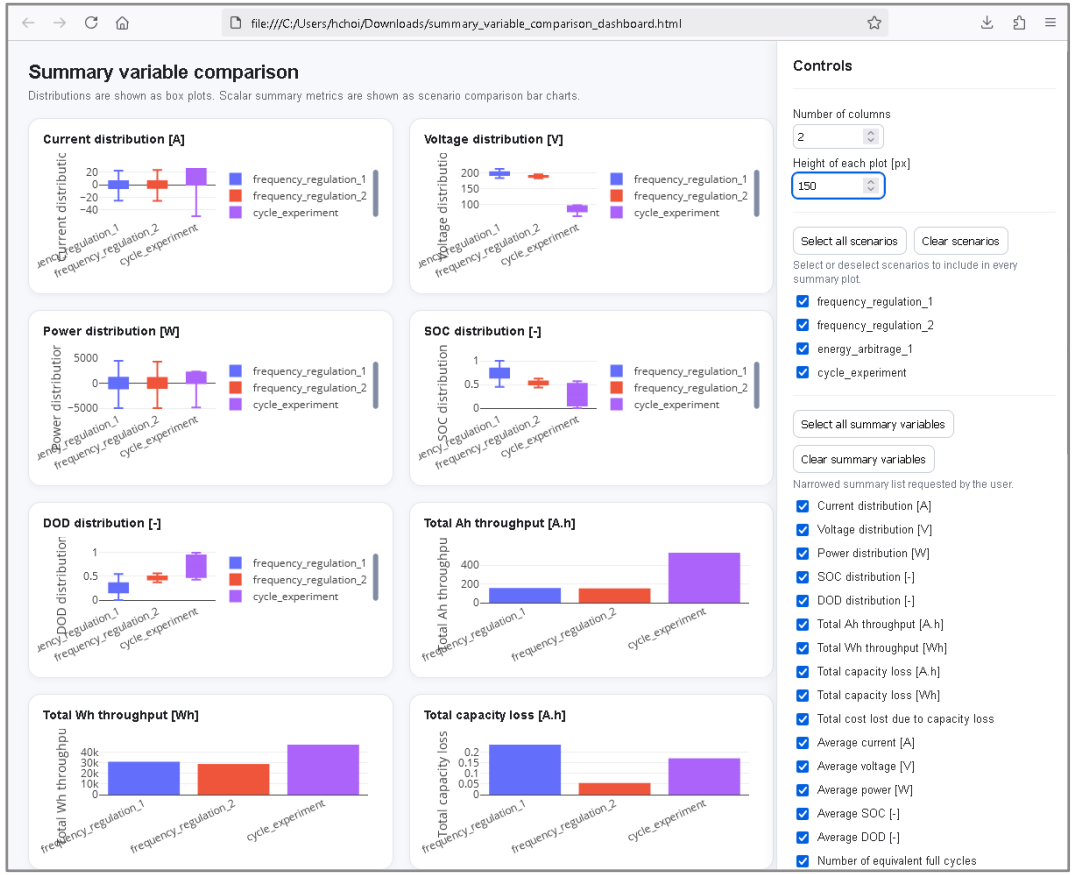

Key summary/statistical plots

QUEST AI TOOL FOR SIMULATION CONFIGURATION SETUP

- **The QuESt AI Agent can help create and modify workflows and scenarios from text-based inputs, which significantly helps non-experts.**

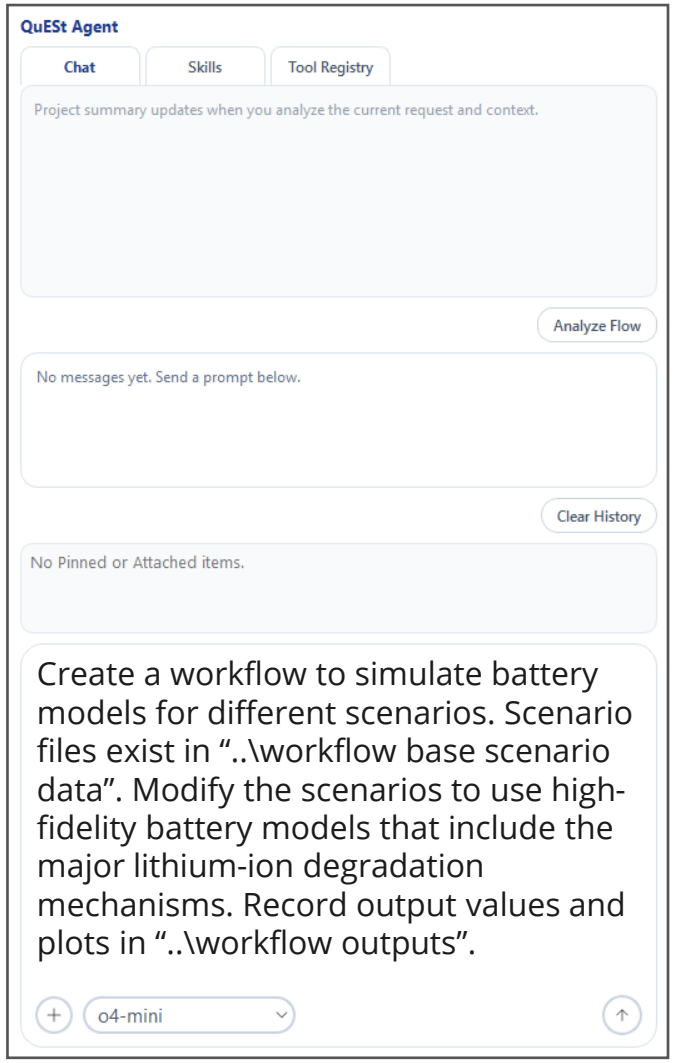

AI-assisted workflow configuration

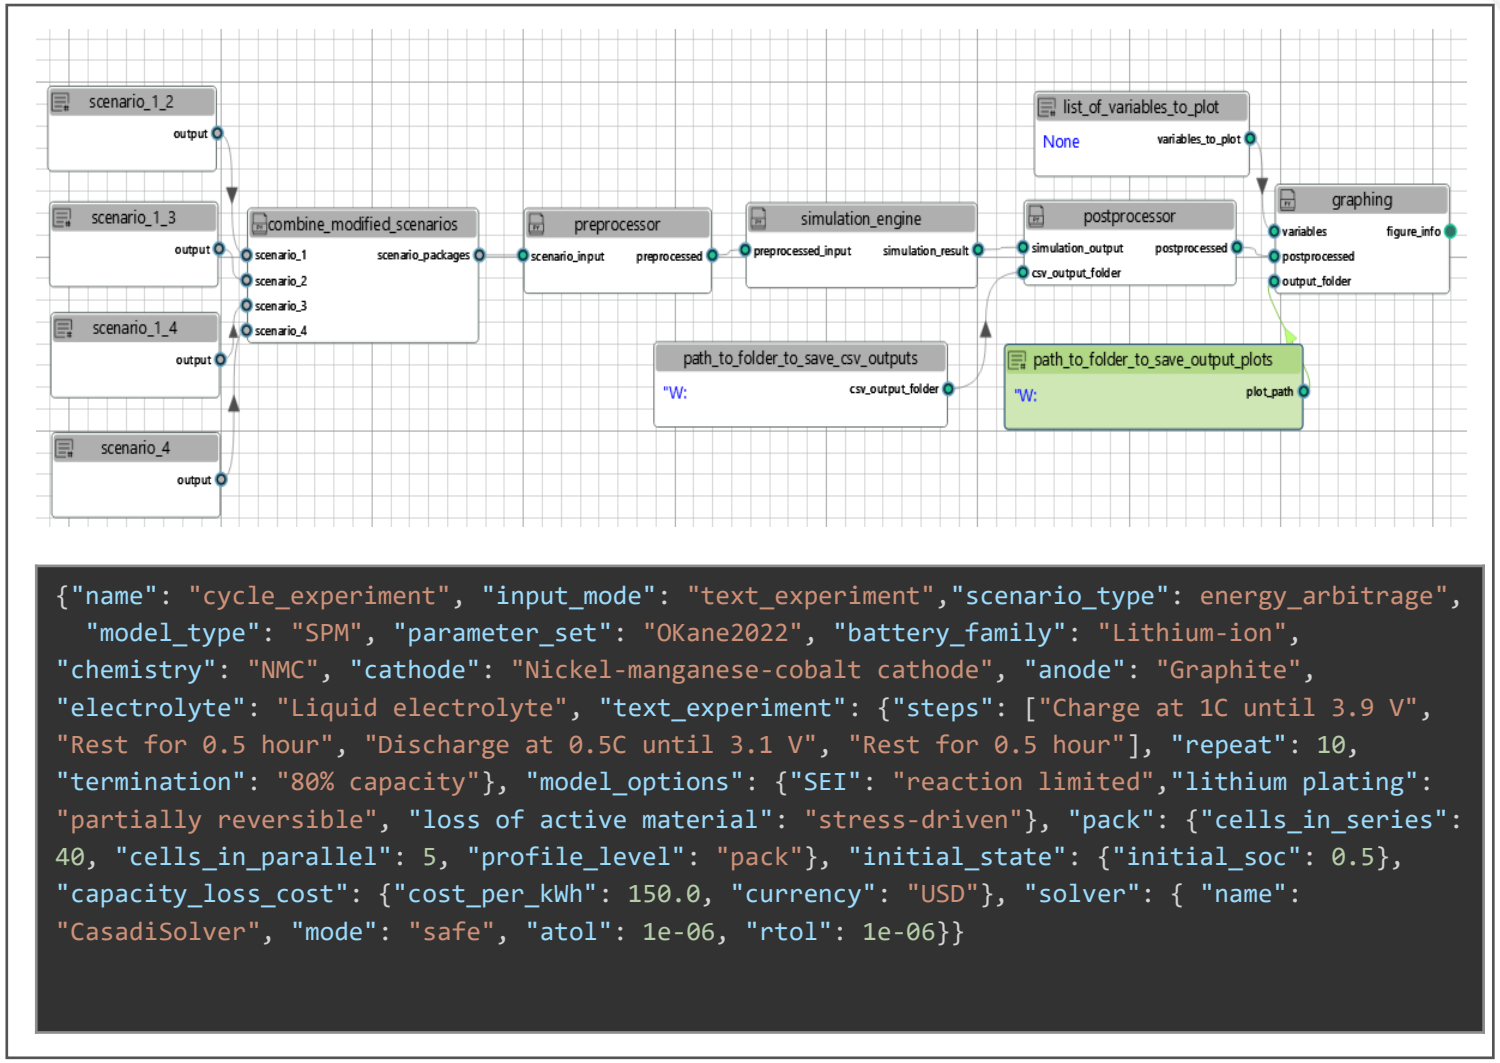

Workflows and scenarios for simulation created by the QuESt AI Agent

ACKNOWLEDGMENT

This material is based upon work supported by *the U.S. Department of Energy, Office of Electricity (OE), Energy Storage Division.*

Sandia National Laboratories

U.S. DEPARTMENT of ENERGY

Contact: Tu Nguyen – tunguy@sandia.gov