AEIC GRID ADVANCEMENT PROGRAM · CHALLENGE BRIEF

Overcoming Interconnection Bottlenecks

Challenge 05. The operating problem, why it is hard, and what the cohort will work.

Working session	Tuesday, December 8, 2026, 2:00 p.m. ET
Cadence	Second Tuesday of every month, 2:00 p.m. ET
Eligibility	AEIC members and any electric utility
Join	https://aeicopscenter.org/grid-advancement
Program	AEIC Center for Operational Excellence

The problem

The interconnection process at most utilities was built for a world where a handful of generators asked to connect each year, and each one got a careful, custom, engineer-led study. That process was not wrong. It was proportionate.

It is not proportionate anymore. Distributed energy resources arrive as a stream now, not as events. Solar, storage, EV charging depots, backup generation somebody wants to sell back to the system, large flexible load looking for a fast path to a substation. The applications keep coming, and the process that clears them is staffed by the same small number of people who can actually run it.

So the queue grows. And a growing queue is not a paperwork problem. It is a planning problem, a capital problem, and increasingly a political one, because the people waiting in it are talking to your commission.

Why this challenge

Because the queue is where the energy transition either happens or does not.

Every clean energy target a state sets, every data center that wants its own generation, every customer who wants storage behind the meter, all of it reaches the utility through the same door. If that door takes eighteen months to open, the utility is the constraint on all of it, whether or not that is fair.

And it is genuinely not fair. The queue does not back up because anybody is slow. It backs up because a serial study process, a distribution model that was never built for this, and a tariff that specifies exactly how each application must be handled all collide with volume nobody staffed for. That collision is worth a room.

Where utilities actually are

The plays in use are mostly the same ones, and mostly reasonable:

Hiring and outsourcing. Add interconnection engineers, or contract the studies out. The fastest lever available, and it hits a ceiling, because only so many distribution engineers can do this work and everybody is hiring them at once.

Screening and triage. Push the small, obviously fine applications through a fast lane so the engineers only touch what needs judgment. Real relief, and it works right up until the fast lane is full too.

Hosting capacity maps. Publish where the system can absorb DER so developers stop applying where the answer was always going to be no. Widely adopted, and only as honest as the model underneath it.

Cluster or batch study. Study applications in groups so one withdrawal does not set off a restudy cascade. Often the right answer, and it changes who pays for what, which turns an engineering fix into a regulatory negotiation.

None of that is wrong. It is what you do when demand goes up an order of magnitude and the process does not.

The hard part

The technology mostly exists. Automated intake, automated screening, power-flow engines that run a study in minutes, hosting capacity analysis that refreshes on a schedule instead of once a year. All real, all working somewhere today. The hard part is not the software. It is these four things.

1. The bottleneck is the engineer, not the paperwork. A portal that accepts applications faster just fills the waiting room faster. The real constraint is the number of people who can look at a marginal application on a

constrained feeder and make a call they can defend. Automate intake without automating judgment and you have moved the pile, not shrunk it. The months are in the study.

2. Automated screening rides on a distribution model that was never built to carry it. Defensible screening needs an as-built model with correct phasing, connectivity, and impedance down to the service transformer. Most distribution GIS was built to track assets and dispatch crews, not to run power flow. The gaps are ones no one ever had a reason to close, and every automated answer inherits them. Precision here is bought with data quality, which is rarely funded and rarely owned, because nothing visible happens on the day you fix it.

3. The queue is a legal process, not a workflow. Study triggers, timelines, cost allocation, and the applicant's rights are written into tariff and state rule. You do not get to skip a step because a model says the answer is obvious. And if an automated screen approves something that later hurts a feeder, the utility owns that, not the tool. That asymmetry is why engineers keep the pen. It is rational, not obstinate.

4. The queue is rarely funded, because it does not show up on the income statement. Interconnection processing is administrative cost that generates no revenue, and the pain of the delay lands on the developer and the customer. So the case for spending real money to go faster is built on somebody else's avoided pain, which is one of the hardest things there is to take to a rate case. That is a structural gap in how the work is compensated, not a lack of anyone's attention.

If you need one sentence to carry into a resourcing conversation: **the months are in the study, not in the paperwork. A faster front door fills the waiting room faster; it does not shrink the queue.**

None of those four barriers is local to you. In some form, all four show up in all ten challenges in this program: the model that was never built for the question being asked of it, the data quality that is rarely funded, the work that falls between functions, and the value that lands on somebody else's balance sheet. If your list looks like this one, you are not behind. You are in the same room as everyone else.

And "we do not actually know where our cycle time goes" is a legitimate answer here. It is a finding, and it is the most useful one you can bring, because almost every real fix in this challenge started with someone pulling the timestamps and being surprised.

What the cohort will work

Not "which platform should we buy." The room is going to work what actually blocks throughput:

- Where do the months actually go? Not where you think, where the timestamps say.
- What has anyone truly automated end to end, and what did they leave with a human on purpose?
- How good does the model have to be before an automated screen is defensible, and what did it cost to get there?
- Who has moved to cluster study, and what did they have to renegotiate to do it?
- What has been tried and did not work? That is the most valuable hour in this program, and it is the one you can only get from peers who have actually tried it.

What to bring

You do not need a solved problem. You need an honest one. Come with:

- Your actual cycle time, and where inside it the time is spent
- What your screening criteria are, and when you last changed them
- The state of your distribution model, honestly assessed
- One thing you tried that did not work

If the honest answer to any of those is "we do not know," bring that. It is a finding, not a failure, and you will not be the only one saying it.

What you get

- This brief, and a short video briefing, before the session
- The Flight Plan, a working pilot scaffold, inside InnovationWorks
- The session itself, with the utilities living this problem
- A retrospective afterward: what the cohort tried and what happened

Join

Tuesday, December 8, 2026, 2:00 p.m. ET. Recruitment runs per challenge, so your interconnection and distribution planning people can come to this one without committing to all ten.

Open to AEIC members and to any electric utility. If you run a utility, you are welcome here.

https://aeicopscenter.org/grid-advancement
Knowledge of one is the knowledge of all.
Sources: this challenge was developed from research of public utility filings. The technology landscape for this challenge is maintained separately inside InnovationWorks.